RORKE DATA®
/ An Avnet Company /

ISO 9001:2008
ISO 13485:2003 Certified

Galaxy DS Firmware Reference Manual

MODELS:

» GX4L-XXXXX Ver 3.85
Galaxy DS Series HDX4 RAID Subsystem firmware ver 3.85

Galaxy DS HDX4 Firmware

7th Generation RAID

With over 10,000 Galaxy units in the field, Rorke Data's award winning RAID products provide the performance, protection, and expansion capabilities for diverse customer environments.

PLEASE READ BEFORE INSTALLATION

» www.rorke.com
/ An Avnet Company /

Rorke Data, An Avnet Company
7626 Golden Triangle Drive, Eden Prairie, MN 55344, USA
» Toll Free 1.800.328.8147 » Phone 1.952.829.0300 » Fax 1.952.829.0988

1.1 Contact Information

Americas

Rorke Data Inc

7626 Golden Triangle Drive

Eden Prairie, MN 55344 USA

Tel: +1-800 328 8147

Fax: +1-952 829 0988

sales@rorke.com

techsupport@rorke.com

http://www.rorke.com

1.2 Copyright 2010

1.2.1 This Edition First Published 2010

All rights reserved. This publication may not be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without the prior written consent of Rorke Data, Inc.

1.2.2 Disclaimer

Rorke Technology makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Rorke Data reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify any person of such revisions or changes. Product specifications are also subject to change without prior notice.

1.2.3 Trademarks

Galaxy and the Galaxy logo are registered trademarks of Rorke Data, Inc.

Solaris and Java are trademarks of Sun Microsystems, Inc.

All other names, brands, products or services are trademarks or registered trademarks of their respective owners.

Table of Contents

1.2.1 This Edition First Published 2010 2
1.2.2 Disclaimer 2
1.2.3 Trademarks 2
2.3.1 Serial Port Settings 9
2.3.2 Activating Windows XP HyperTerminal 9
2.3.3 The Firmware Interface at a Glance 12
2.4.1 Configuring the RS-232 Port 13
2.4.2 Configuring Terminal Emulation 14
2.4.3 Configuring the Baud Rate 14
2.4.4 Configuring Internet Protocol <TCP/IP> 14
2.4.5 Viewing the Link Status 15
2.4.6 Configuring the IP Address 15
2.4.7 Configuring the Prefix Length field 17
2.4.8 Disabling Network Protocols 18
2.5.1 Connecting to the Ethernet Port 19
2.5.2 Configuring the Controller 20
2.5.3 Connecting through Telnet 21
2.5.4 Establishing Secure Link over SSH 22
3.1.1 The Initial Screen 24
3.1.2 About Logical Drives 24
3.1.3 Logical Drive Status 25
3.1.4 Logical Volume Status 26
3.1.5 Physical Drive Status 27
3.1.6 Channel Status 28
3.1.7 Viewing Controller Voltage and Temperature 29
3.1.8 Viewing and Editing Event Logs 30
3.2.1 The Initial Screen 31
3.2.2 Main Menu 32
3.2.3 Notes on Logical Drives 33
3.2.4 Viewing Logical Drive Status 34
3.2.5 Viewing Logical Volume Status 36
3.2.6 Viewing Physical Drive Status 37
3.2.7 About EXILED Drives 40
3.2.8 Viewing Channel Status 41
3.2.9 Viewing Controller Voltage and Temperature 43
3.2.10 Viewing Event Logs on Screen 44
4.1.1 Deciding the Stripe Size 46
4.1.2 Enabling Write-Back Cache 47
4.1.3 Enabling Write-Back 48
4.1.4 Trigging Events 49
4.1.5 Flushing Cache Periodically 51
4.2.1 Notes on Channel Mode Settings 52
4.2.2 Configuring Channel ID 53
4.2.3 Adding a Host ID 56
4.2.4 Deleting an ID 59
4.2.5 Selecting the Data Rate (Host Channel Bus) 60
4.2.6 Selecting the Data Rate (Drive Channel) 61
4.3.1 Selecting the Time Zone 63
4.3.2 Setting the Date and Time 64
4.5.1 About Auto-assignment of a Global Spare 66
4.5.2 Auto-Assigning a Global Spare 68
4.5.3 About Enclosure Spare 69
4.5.4 Assigning an Enclosure Spare 70
4.7.1 Muting Beeper Sound 72
4.7.2 Changing the Password 73
4.7.3 Resetting the Controller 75
4.7.4 Shutting Down the Controller 76
4.7.5 Saving NVRAM to Disks 77
4.7.6 Restoring NVRAM from Disks 78
4.7.7 Clearing Core Dump 79
4.7.8 Adjusting the LCD Contrast 80
4.8.1 Changing the Controller Name 81

4.8.2 Showing the Controller Name 82
4.8.3 Setting Password Validation Timeout 82
4.8.4 Setting a Unique Controller Identifier 83
5.1.1 List of Key Differences 86
5.1.2 87
Storage Components 87
5.1.3 Data Services 88
5.1.4 Limitations 88
5.2.1 Typical HDX4 Deployment 89
6.3.1 Creating a Logical Drive 91
6.3.2 Choosing Member Drives 92
6.3.3 Setting Maximum Drive Capacity 92
6.3.4 Assigning a Spare Drive 93
6.3.5 Viewing Reserved Disk Space 93
6.3.6 Setting Write Policy 94
6.3.7 Setting Initialization Mode 94
6.3.8 Setting Stripe Size 95
6.3.9 Initializing a Logical Drive 98
6.3.10 Naming a Logical Drive 99
6.3.11 Deleting a Logical Drive 99
6.3.12 Deleting the Partition of a Logical Drive 100
6.4.1 Creating a Logical Volume 102
6.4.2 Setting the Initialization Mode 102
6.4.3 Setting the Write Policy 103
6.4.4 Assigning a Logical Volume (Dual-active Controllers) 104
6.4.5 Partitioning a Logical Volume 105
6.4.6 Mapping Logical Partitions Drive to Host LUN 106
6.4.7 Deleting Host LUNs 108
6.5.1 Adding a Local Spare Drive 109
6.5.2 Adding a Global Spare Drive 109
6.5.3 Adding an Enclosure Spare Drive 110
6.5.4 Deleting Spare Drive (Global / Local/Enclosure Spare Drive) 111
7.2.1 Creating a Logical Drive 114
7.2.2 Setting Maximum Drive Capacity 118
7.2.3 Assigning Spare Drives 118
7.2.4 Changing Logical Drive Assignments 119
7.2.5 Changing Write Policy 120
7.2.6 Setting the Initialization Mode 120
7.2.7 Setting the Stripe Size 121
7.2.8 Setting the Power Saving Mode 122
7.2.9 Editing Logical Drives 124
7.2.10 Deleting a Logical Drive 125
7.2.11 Naming a Logical Drive 125
7.2.12 Expanding a Logical Drive or a Logical Volume 126
7.3.1 Requirements for Migrating a RAID5 Array 127
7.3.2 Migration Methods 128
7.3.3 Migration: Exemplary Procedure 129
7.4.1 Creating a Logical Volume (Required) 131
7.4.2 Notes on Partitions in Galaxy HDX4 Series 133
7.4.3 Creating a Partition from Logical Volume (Required) 134
7.4.4 Deleting Partitions 136
7.4.5 Examining Valid Connectivity 137
7.4.6 Managing Host Adapter Ports 138
7.4.7 Notes on Mapping a Partition to Host LUN 140
7.4.8 Mapping a Partition to a Host LUN 141
7.4.9 Deleting Host LUNs 143
7.4.10 Expanding a Logical Volume 143
7.5.1 Adding a Local Spare Drive 145
7.5.2 Adding a Global Spare Drive 145
7.5.3 Adding an Enclosure Spare Drive 146
7.5.4 Deleting Spare Drive (Global/Local/Enclosure Spare Drive) 146
8.1.1 Channel IDs - Host Channel 147
8.1.2 Adding an ID (Slot A / Slot B Controller ID) 148
8.1.3 Deleting an ID 149
8.1.4 Setting Data Rate (Channel Bus) 150
8.1.5 Viewing Channel Host ID/WWN 151
8.1.6 Viewing Device Port Name List (WWPN) 153

8.1.7 Adding Host – ID/WWN Label Declaration 153
8.2.1 About Loop Only 154
8.2.2 About Point-to-point 154
8.2.3 Setting Controller Unique Identifier 155
9.1.1 iSCSI IP SAN – 1 Topology 157
9.1.2 iSCSI IP SAN – 2 Topology 158
9.1.3 iSCSI IP SAN – 3 Topology 159
9.1.4 iSCSI IP SAN – 4 Topology 160
9.2.1 Setting Switch Trunk Port 162
9.2.2 Notes on Trunking Conditions 164
9.2.3 Configuring Trunk 166
9.3.1 Grouping VS Trunking 168
9.3.2 Configuring Group 169
9.3.3 LUN Presentation with and without Grouping 170
9.3.4 Channels Automatically Divided into A and B Sub-groups 171
9.3.5 LUN Presentation on Multiple Data Paths 172
9.4.1 IP Addresses to the iSCSI Host Ports 173
9.4.2 Creating Host Channel IDs 174
9.5.1 Creating an iSCSI Initiator List 176
9.5.2 Configuring Initiator (Using Microsoft Software Initiator) 177
9.5.3 About IQN Name 178
9.5.4 Sample IQN Procedure 179
9.6.1 Configuring CHAP on RAID 185
9.6.2 Configuring CHAP on the Initiator 187
9.7.1 iSNS Overview 192
9.7.2 iSNS Configuration Sample and Flowchart 193
9.7.3 iSNS Configuration (RAID) 193
9.7.4 iSNS Configuration (PC) 194
9.8.1 How does it work? 196
9.8.2 Configuring SLP 198
10.1.1 Maximum Concurrent Host LUN Connection (“Nexus” in SCSI) 202
10.1.2 Number of Tags Reserved for Each Host-LUN Connection 204
10.1.3 Maximum Queued I/O Count 204
10.1.4 LUNs per Host ID 205
10.1.5 LUN Applicability 206
10.1.6 Peripheral Device Type 206
10.1.7 In-band Management Access 206
10.1.8 Peripheral Device Type Parameters for Various Operating Systems 207
10.1.9 Cylinder/Head/Sector Mapping 208
10.2.1 Disk Access Delay Time 210
10.2.2 Drive I/O Timeout 211
10.2.3 Maximum Tag Count: Tag Command Queuing (TCQ) and Native Command Queuing (NCQ) Support 211
10.2.4 Drive Delayed Write 212
10.2.5 Power Saving 213
11.1.1 RAID Enclosure Devices 215
11.1.2 Devices within the Expansion Enclosure 216
11.1.3 Verifying Disk Drive Failure in a Multi-enclosure Application 217
11.2.1 Event Triggered Operations 218
11.2.2 Operation Theory 219
11.2.3 Auto Shutdown on Elevated Temperature 220
11.2.4 Voltage and Temperature Self-monitoring 221
11.2.5 Changing Monitoring Thresholds 221
12.2.1 Auto Rebuild on Drive Swap Check Time 225
12.2.2 Auto-Assign Global Spare Drive 226
12.3.1 Task Scheduler 227
12.3.2 Setting Task Scheduler 227
12.5.1 Overwriting Inconsistent Parity 234
12.5.2 Generating Check Parity Error Event 234
12.6.1 Rebuild Priority 235
12.6.2 Verification on Writes 235
13.1.1 What is RAID Expansion and how does it work? 237
13.1.2 Notes on Expansion 237
13.1.3 Expand Logical Drive: Re-striping 238
13.2.1 Overview 239
13.2.2 Add Drive Procedure 239
13.3.1 Overview 242

13.3.2	Copy and Replace Procedure	242
13.4.1	Expanding Logical Drives	244
13.4.2	Expanding Logical Volumes	245
13.5.1	Prerequisites	246
13.5.2	Step 1: Expanding the Logical Drives	247
13.5.3	Step 2: Expanding the Logical Volume	248
13.5.4	Step 3: Expanding the Partition	248
13.5.5	Step 4: Expand the Original Logical Volume in Computer Management Utility	249
14.1.1	Replacing After Clone	251
14.1.2	Perpetual Clone	253
14.2.1	Introduction	255
14.2.2	Galaxy's Implementations with S.M.A.R.T.	256
14.3.1	Enabling the S.M.A.R.T. Feature	257
14.3.2	Using S.M.A.R.T. Functions	258
15.1.1	Fewer Streams: Read-ahead Performance	260
15.1.2	Multi-Streaming: Simultaneous Access Performance	262
15.2.1	Response Time in Read Scenarios	263
15.2.2	Maximum Drive Response Time in Write Scenarios	264
15.2.3	Other Concerns	265
16.1.1	Concerns	266
16.1.2	Communications Channels	267
16.1.3	Out-of-Band Configuration Access	268
16.1.4	Limitations	268
16.1.5	Configurable Parameters	269
16.2.1	General Firmware Configuration Procedures	269
16.2.2	Setting Controller Unique ID (Optional)	271
16.2.3	Creating Controller A and Controller B IDs	272
16.2.4	Logical Volume Assignments (Dual-Controller Configuration)	273
16.2.5	Mapping a Logical Volume to Host LUNs	277
16.3.1	What will happen when one of the controllers fails?	279
16.3.2	When and how is the failed controller replaced?	280
16.3.3	How Do I Resolve Conflict in Assigning the Primary Controller?	282
16.4.1	RCC (Redundant Controller Communications Channel) Status	283
16.4.2	Adaptive Write Policy	284
16.4.3	Adaptation for the Redundant Controller Operation	285
16.4.4	Cache Synchronization on Write-Through	285
16.5.1	The Inter-Controller Relationship	286
16.5.2	Rules for Grouping Hard Drives and LUN Mapping	286
16.5.3	Host LUN Mapping: Design Concerns	289
16.5.4	Mapping for Fault-tolerant Links	289
16.5.5	Mapping Using the Cross-controller Mapping	292
16.5.6	Fault Tolerance	293
16.5.7	Fault Tolerance Procedures	294
16.5.8	Controller Failure	295
16.6.1	Design Concerns	296
16.6.2	Simple DAS without Hub (Cross-controller Mapping Method)	297
16.6.3	SAN with FC Switches	299
16.6.4	Multi-pathing with Clustered Servers (Cross-controller Mapping Method)	301
18.1.1	Sample Flowchart	315
18.1.2	Note for Redundant Controller Firmware Upgrade	315
18.3.1	Upgrading Both Boot Record and Firmware Binaries	318

About This Manual

This manual describes Firmware for the Galaxy Data Services [DS] RAID Series.

The Data Service features include: Snapshot, Volume Copy, Volume Mirror, Thin Provision (from firmware version 386), and a scheduler tool for these features. Refer to the Galaxy Array Manager manual to apply those features.

This includes only the Galaxy HDX4 RAID model. Earlier models of Galaxy HDX3, HDX2, and HDX RAIDs do not have these features nor can these features be applied to these earlier models.

Revision History

Version	Description	Date
1.2	Changed document format Updated content	November 2010

Applicable Firmware Version

This manual is applicable to FW 385 or later.

2 Establishing Connections

This chapter describes how to establish the management access to your RAID system. The main topics include the following:

- RS-232C Serial Port
- Communication Parameters
- Out-of-Band via Ethernet
- Telnet Connection
- Secure Link over SSH

2.3 Working with RS-232C Serial Port

The Galaxy Data Services [DS] series firmware can be configured in a text user interface and can be accessed through a terminal emulator application:

- Windows XP or before: The HyperTerminal program pre-installed in the OS
- Other OS: A terminal emulator application such as the VT100 series emulator.

To access the firmware, you may use the RS-232C interface provided with the Galaxy storage subsystems.

- For dual-controller subsystems, use the serial Y-cable included in the package.
- For single-controller subsystems, the serial cable is user-provided. A standard DB9 serial cable can be used.

NOTE

The connection is straight-line. No null modem adapter is required.

For computers without an RS-232C port, you may use a USB-to-DB9 adapter.

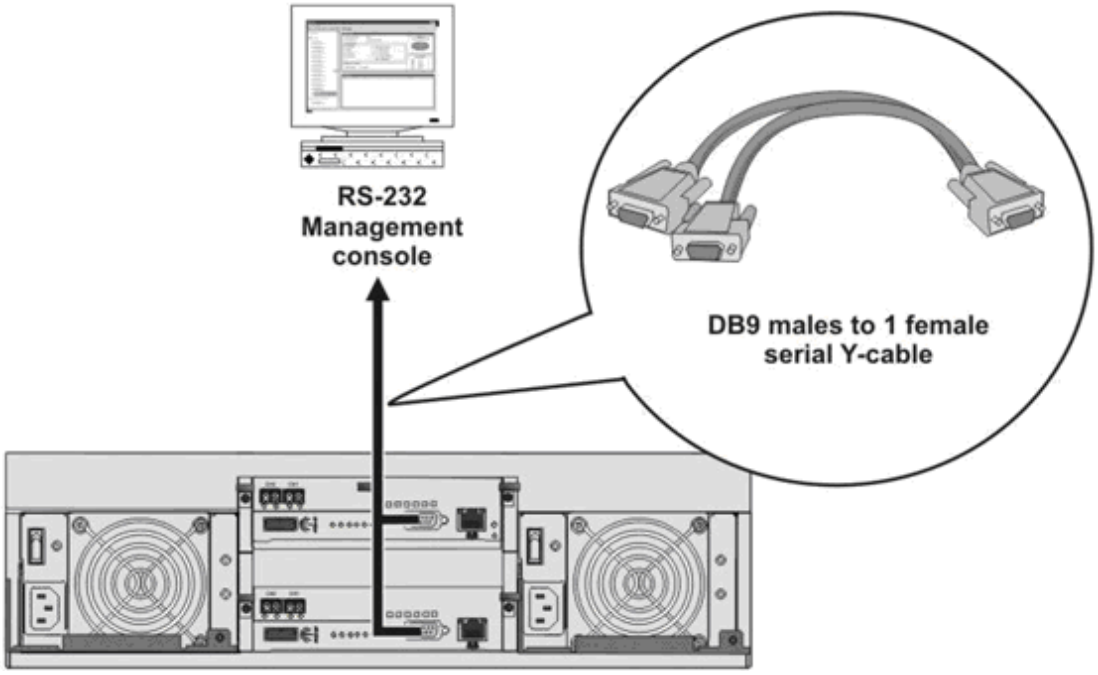

RS-232C Connection (Dual-Controller Model)

2.3.1 Serial Port Settings

We recommend the following configurations for the RS-232C serial port.

Baud Rate	38400
Data Bit	8
Parity	None
Stop Bit	1
Flow Control	Hardware
COM Port	COM1

2.3.2 Activating Windows XP HyperTerminal

1. Select *Start* > *Accessories* > *Communications* > *Hyper Terminal*.
2. Enter the country, area code, and the connection name.

3. Select the COM port.

4. Set the parameters (you may follow these recommendations).

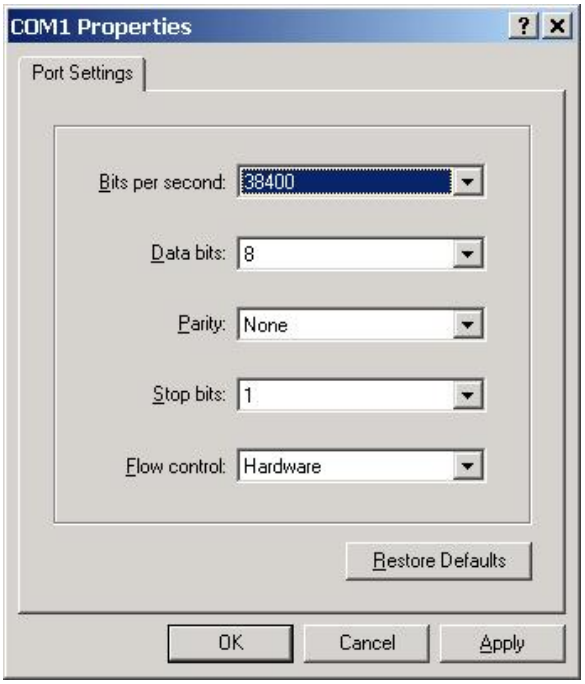

5. If pop-up messages appear, press the ESC key to clear them.

6. The Galaxy HDX4 firmware screen will appear and starts the initial test.

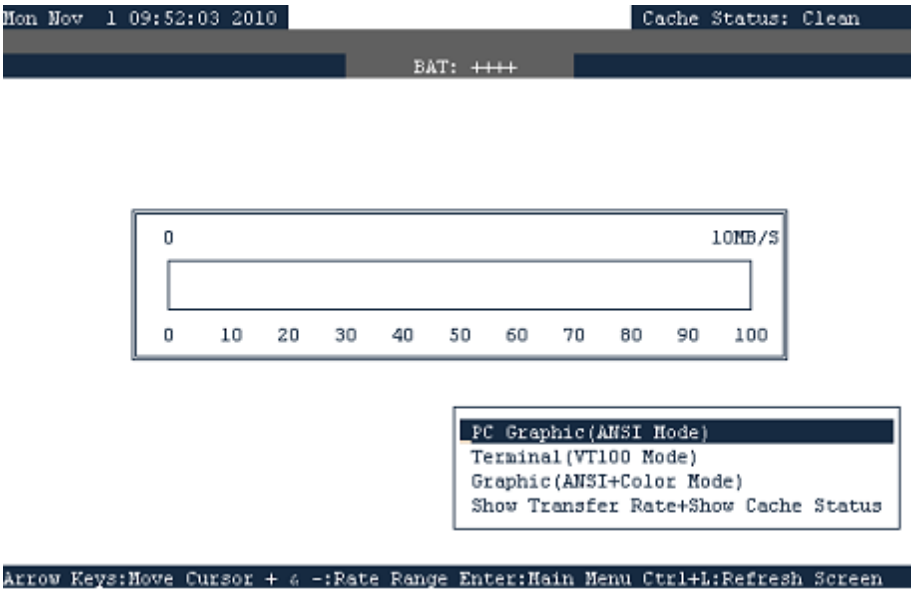
Mon Nov 1 09:52:03 2010
Cache Status: Clean
BAT: ++++
0
10MB/S
0 10 20 30 40 50 60 70 80 90 100
PC Graphic(ANSI Mode)
Terminal(VT100 Mode)
Graphic(ANSI+Color Mode)
Show Transfer Rate+Show Cache Status
Arrow Keys:Move Cursor + & -:Rate Range Enter:Main Menu Ctrl+L:Refresh Screen

7. The firmware main menu will appear.

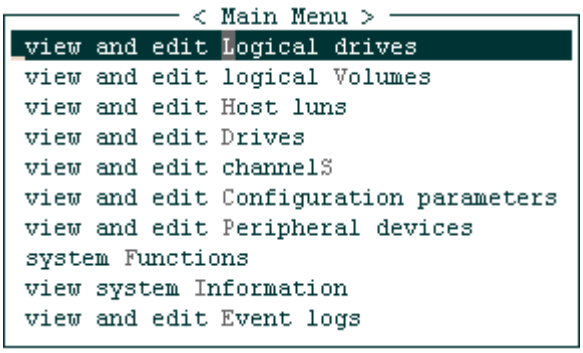
< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channelS
view and edit Configuration parameters
view and edit Peripheral devices
system Functions
view system Information
view and edit Event logs

8. To refresh the screen status, press Ctrl+L keys. Now you are ready to go.

2.3.3 The Firmware Interface at a Glance

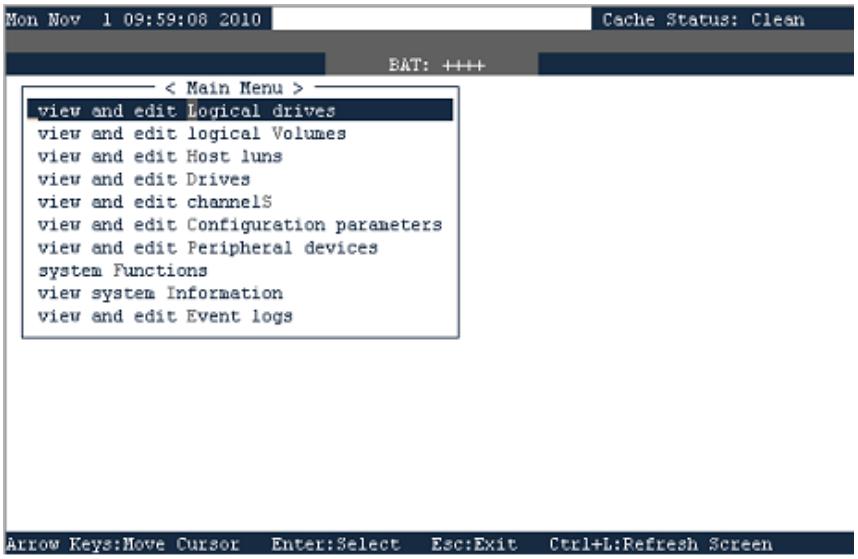

Date	Shows the current date and time. You may reconfigure it.
Cache Status	Shows the storage subsystem's cache memory usage.
CBM Name	Shows the name of the controller.
Menu	Shows the current menu options.
Keys	Shows the control options.

Use the following keys to navigate and operate in this interface.

[Arrow Keys]	Selects menu options.
[Enter]	Executes the selected option or enters submenus.
[Esc]	Cancels an option or returns to the previous menu.

[Ctrl]+[L] Refreshes the screen information

The initial screen appears when the controller finishes its self-test and is properly initialized. Use Up/Down arrow keys to select terminal emulation mode, then press [ENTER] to enter the Main Menu.
Choose a functional item from the Main Menu to begin configuring your RAID.

2.4 Configuring Parameters

Go to: View and Edit Configuration Parameters > Communication Parameters

The Communication Parameters is the first sub-menu under the “View and Edit Configuration Parameters” menu. In addition to the baud rate and terminal emulation options which have been discussed earlier, the sub-menu contains other options to prepare your management session using an Ethernet connection.

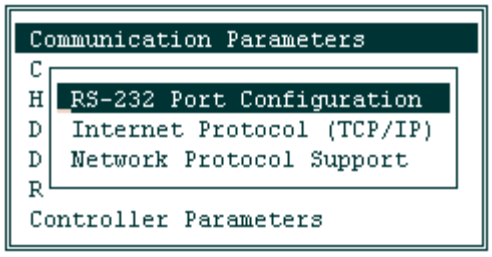

2.4.1 Configuring the RS-232 Port

Go to: View and Edit Configuration Parameters > Communication Parameters > RS-232 Port Configuration

The “RS-232 Port Configuration” provides access to change the serial port operating parameters. Each COM port (COM1) selection menu features two communication parameters: “Baud Rate” and “Terminal Emulation.”

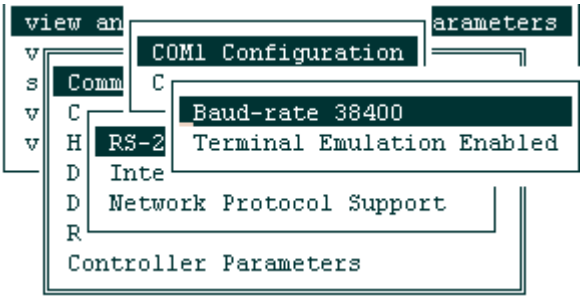

NOTE

On HDX4 series models, the COM2 serial port is cancelled.

2.4.2 Configuring Terminal Emulation

Go to: View and Edit Configuration Parameters > Communication Parameters > Terminal Emulation

The Terminal Emulation setting on the COM1 port is enabled by default. Usually there is no need to change this setting.

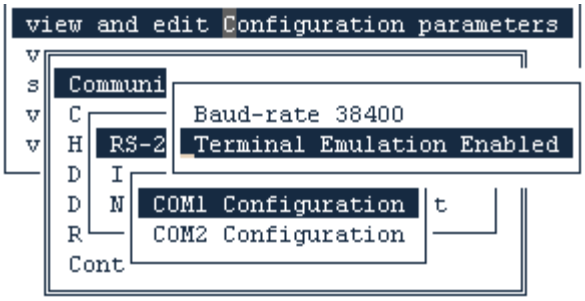

2.4.3 Configuring the Baud Rate

Go to: View and Edit Configuration Parameters > Communication Parameters > RS-232 Port Configuration > Baud-rate

Available options will be displayed in a pull-down menu. Select by pressing [ENTER] and press ESC several times to return to the previous configuration screen.

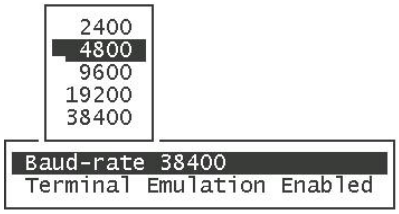

2.4.4 Configuring Internet Protocol <TCP/IP>

Go to: View and Edit Configuration Parameters > Communication Parameters > Internet Protocol

The Internet Protocol menu allows you to prepare the management access through the system's RJ-45, 10/100BaseT Ethernet port.

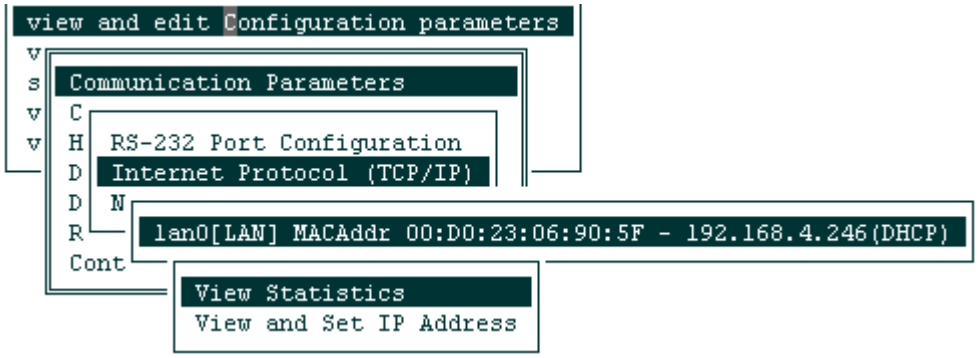

To access the configuration options, press [ENTER] on “Internet Protocol <TCP/IP>” to display the information of Ethernet port. Press [ENTER] on the chip information to display the “View Statistics” and the “Set IP Address” options.

2.4.5 Viewing the Link Status

Go to: View and Edit Configuration Parameters > Communication Parameters > Internet Protocol > (TCP/IP) > View Statistics

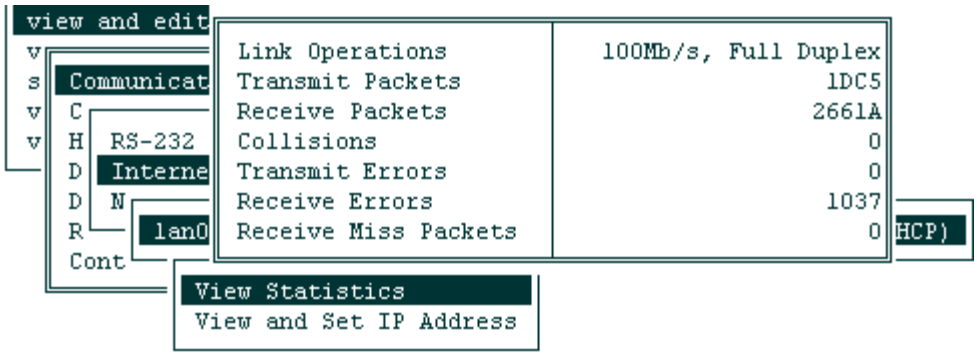

This window displays the current Ethernet link status.

2.4.6 Configuring the IP Address

Go to: View and Edit Configuration Parameters > Communication Parameters > Internet Protocol > (TCP/IP) > View and Set Up IP Address

Provide a valid IP address for your subsystem/controller’s Ethernet port. Consult your network administrator for a static IP address and the associated NetMask and Gateway values. You may also key in “DHCP” if your local network supports automatic IP configuration.

NOTE

The IP default is “DHCP client.” However, if DHCP server can not be found within several seconds, a default IP address “10.10.1.1” will be loaded.

One drawback of using DHCP is that if cable disconnection or other unpredictable network faults occur, your Ethernet port may be assigned with a different IP. This may cause problems for the management sessions using Galaxy Array Manager. You may not be able to receive important event messages before you access the array by entering a new IP address.

It may take several minutes to obtain an IP address from the DHCP server.

Internet Protocol Version 6 (IPv6) is supported on Galaxy HDX4 RAIDs.

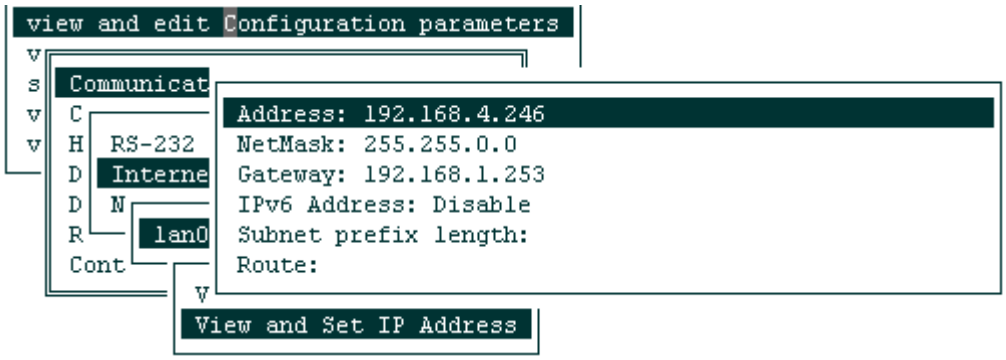

Since IPv6 comes with a more autonomous support for automatic addressing, automatic network configuration is applied in most deployments. An automatic local name resolution is available with or without a local Domain Name Server (DNS).

To assign an IPV6 address automatically, key in “AUTO” in the IPv6 address field.

IPv6 addresses can be acquired through the following ways:

- A link-local address is automatically configured by entering AUTO in the IPv6 address field. With a point-to-point connection without router, addresses will be generated using port MAC addresses starting with “fe80::.” Link-locals are addresses within the same subnet.
- If addresses are automatically acquired, the “Subnet prefix length” and the “Route” fields can be left blank.
- A DHCPv6 server, if present in the network, will be automatically queried for an IPv6 address.
- If an IPv6 router is present, you can Key in AUTO in the Route field and let a

router’s advertisement mechanism determine network addresses.

- You can also manually enter IPv6 addresses by generating the last 64 hexadecimal bits from the 48-bit MAC addresses of Ethernet ports in EUI-64 format, and then use the combination of fe08 prefix and prefix length to signify a subnet.
- A sample process is shown below: 1. Insert FFFE between company ID and node ID, as the fourth and fifth octets (16 bits). 2. Set the Universal/Local (U/L) bit, the 7th of the first octet, to a value of 0 or 1. “0” indicates a locally administered identity, while “1” indicates a globally unique IPv6 interface ID.

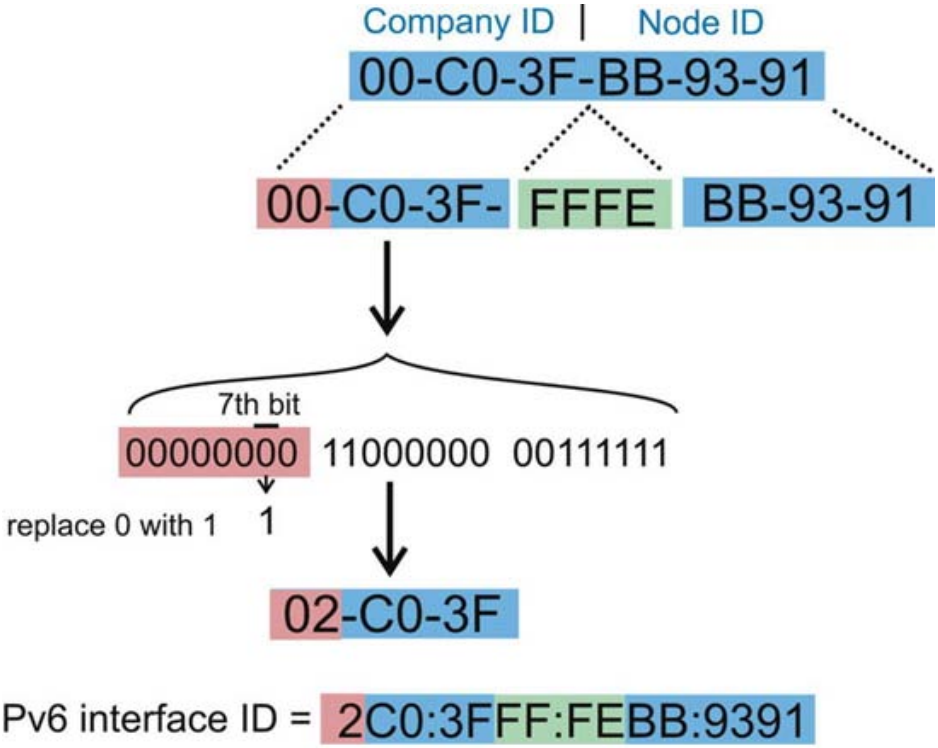

Galaxy supports a variety of IPv6 mechanisms including Neighbor Unreachability Detection, stateful and stateless address auto configuraion, ICMPv6, Aggregatable Global Unicast Address, Neighbor Discovery, etc.

2.4.7 Configuring the Prefix Length field

The prefix length is part of the manual setting. An IPv6 network is a contiguous group of IPv6 addresses. The size of this field must be a power of 2. The Prefix Length designates the number of bits for the first 64 bits of the Ipv6 addresses, which are identical for all hosts in a given network, are called the network's address prefix.

Such consecutive bits in IPv6 addresses are written using the same notation previously developed for IPv4 Classless Inter-Domain Routing (CIDR). CIDR

notation designates a leading set of bits by appending the size (in decimal) of that bit block (prefix) to the address, separated by a forward slash character (/), e.g., 2001:0db8:1034::5678:90AB:CDEF:5432/48. (In firmware screen, slash is not necessary. The prefix number is entered in the length field.)

The architecture of IPv6 address is shown below:

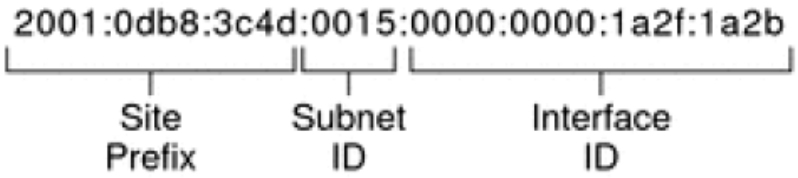

The first 48 bits contain the site prefix, while the next 16 bits provide subnet information. An IPv6 address prefix is a combination of an IPv6 prefix (address) and a prefix length. The prefix takes the form of “ipv6-prefix/prefix-length” and represents a block of address space (or a network). The ipv6-prefix variable follows general IPv6 addressing rules (see RFC 2373 for details).

For example, an IPv6 network can be denoted by the first address in the network and the number of bits of the prefix, such as 2001:0db8:1234::/48. With the /48 prefix, the network starts at address 2001:0db8:1234:0000:0000:0000:0000:0000 and ends at 2001:0db8:1234:ffff:ffff:ffff:ffff:ffff.

Individual addresses are often also written in CIDR notation to indicate the routing behavior of the network they belong to. For example, the address 2001:db8:a::123/128 indicates a single interface route for this address, whereas 2001:db8:a::123/32 may indicate a different routing environment.

IPv6 Prefix	Description
2001:410:0:1::45FF/128	A subnet with only one Ipv6 address
2001:410:0:1::/64	A subnet that contains 264 nodes. Often the default prefix length for a subnet.
2001:410:0::/48	A subnet that contains 216 nodes. Often the default prefix length for a site.

2.4.8 Disabling Network Protocols

Go to: View and Edit Configuration Parameters > Communication Parameters > Network Protocol Support

You may disable one or more of the network protocols to lower the risk of network

security problems.

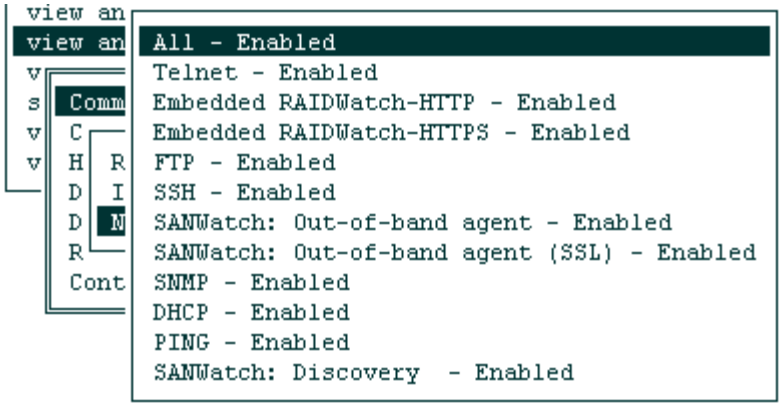

2.5 Connecting Out-of-Band via Ethernet

2.5.1 Connecting to the Ethernet Port

Use a LAN cable to connect the Ethernet port(s) on the system's RAID controller unit(s). Connect the cables between the system's Ethernet port and an Ethernet port on your local network.

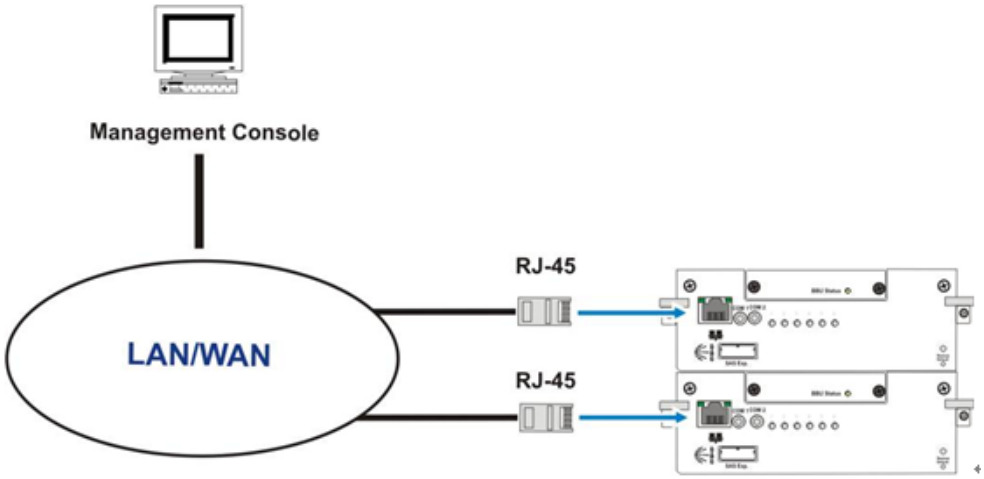

For dual-controller subsystems, connect the Ethernet interfaces from both controllers to your Ethernet network. The Ethernet port on the Secondary controller stays idle and becomes active in the event of Primary controller failure. The Ethernet port IP on a Primary's Ethernet port will be inherited by the Secondary controller during a controller failover process.

NOTE

Due to the high risk of network attack in today's Internet, the Galaxy HDX4's 10/100BaseT management port should always be connected to a local network with

reasonable protection, such as firewall, router, and ISA VPN between trusted and untrusted networks. It is not recommended to directly assign an unprotected public IP to a system's management port.

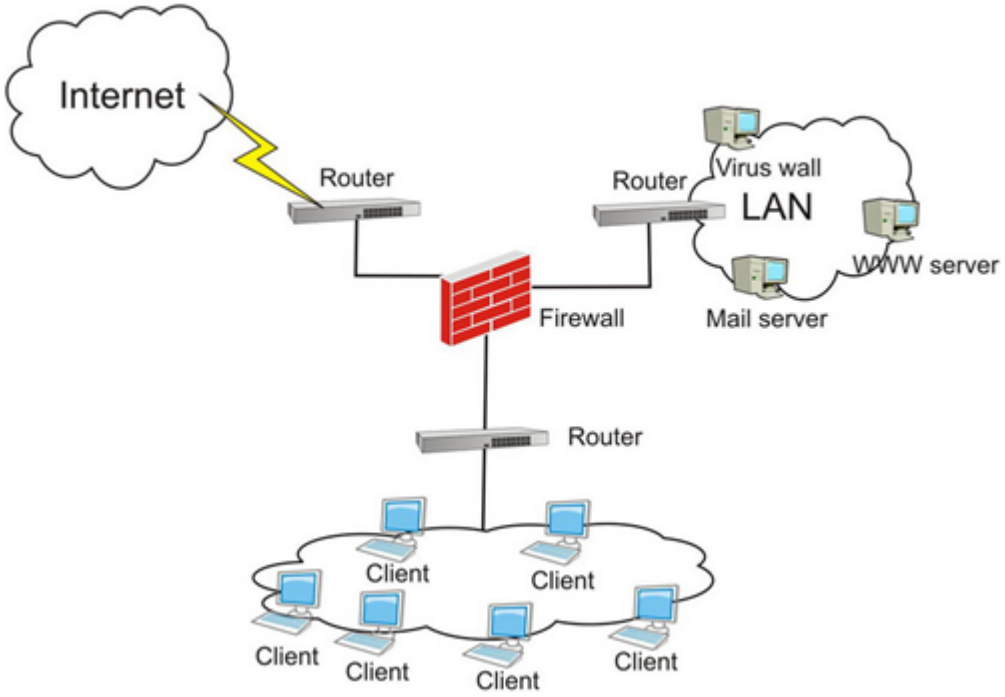

2.5.2 Configuring the Controller

To prepare the subsystem/controller for Ethernet connection:

1. Connect the subsystem's serial port to a PC running a VT 100 terminal emulation program or a VT-100-compatible terminal using the included serial cables.

2. Make sure the included null modem is already attached to the enclosure serial port or the management computer's COM port. The null modem converts the serial signals for connecting to a standard PC serial interface.

3. Go to: View and Edit Configuration Parameters > Communication Parameters > Internet Protocol > (hardware) > View and Set IP Address.

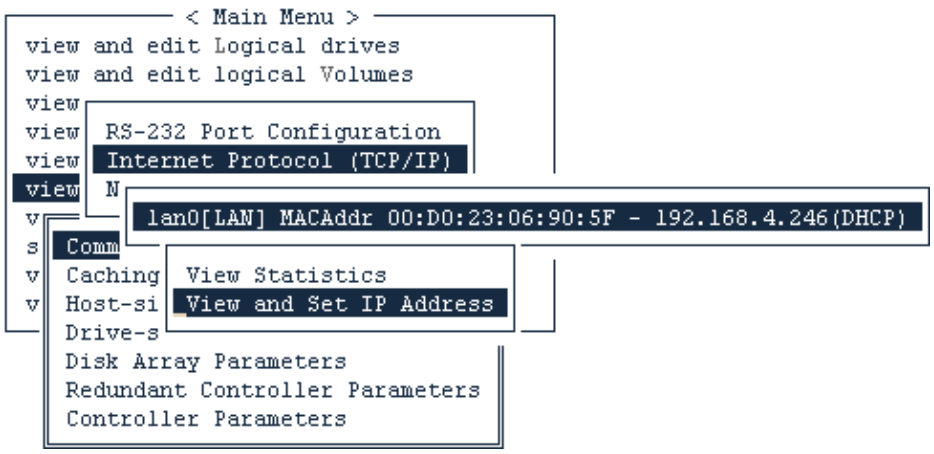

4. You may also use an auto discovery protocol such as DHCP. Simply key in “DHCP” in the IP address field.

5. Provide the IP address, NetMask, and Gateway values accordingly.

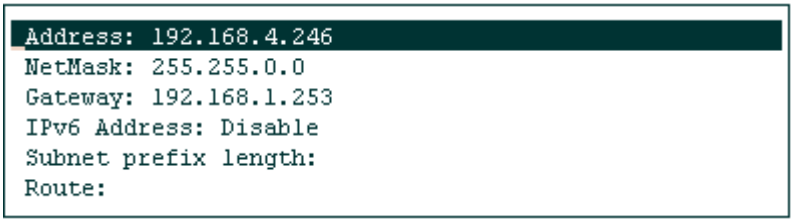

6. PING the IP address from your management computer to make sure the link is up and running.

2.5.3 Connecting through Telnet

1. Use an Ethernet cable with RJ-45 phone jacks to connect the Ethernet port on the subsystem/controller module.

2. Connect the other end of the Ethernet cable to your local area network. An IP address should be acquired for the subsystem’s Ethernet port. The subsystem firmware also supports automatic client configuration such as DHCP.

3. Consult your network administrator for an IP address that will be assigned to the subsystem/controller Ethernet port.

4. Select "View and Edit Configuration Parameters" from the Main Menu on the terminal screen. Select "Communication Parameters" -> "Internet Protocol (TCP/IP)" -> press ENTER on the chip hardware address -> and then select "Set IP Address."

5. Provide the IP address, NetMask, and Gateway values accordingly.

6. PING the IP address from your management computer to make sure the link is valid.

7. Open a command prompt window and key in “telnet xxx.xxx.xx.xxx (IP address)” to access the embedded firmware utility. The default port number is 23.

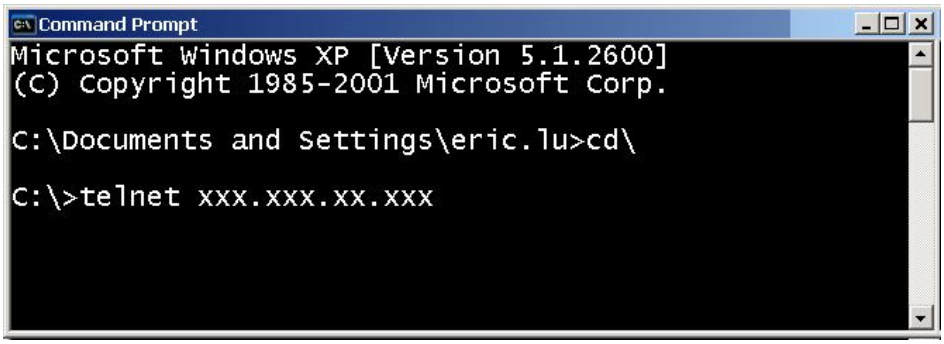

NOTE

When using Telnet, ALWAYS log out using the proper method (ESC key). NEVER finish the session in other ways, such as Closing the Command Prompt or applications such as PuTTY. Doing so might lead to system error.

2.5.4 Establishing Secure Link over SSH

Firmware supports remote management over the network connection and the security under SSH (Secure Shell) protection. SSH is widely used for its ability to provide strong authentication and secure communications over insecure channels. SSH secure access can also be found as an option in the Galaxy Array Manager software.

SSH is more readily supported by Linux- or Unix-based systems. The support for SSH on Microsoft Windows platforms can be limited. For making SSH link using Windows, there are SSH tools such as the “PuTTY” shareware.

To make an SSH link, use “root” as the default user name. If you have configured a controller name and password for your Galaxy HDX4 system, use them as your login name and password.

If a shareware is used, it may be necessary to configure the display options, e.g., the

“Character set translation on received data” and “font type” setting in order for the terminal screen to be correctly displayed. The appearance settings may vary on different SSH tools. The default port number is 22.

3 Screen Messages

3.1 LCD Screen Messages

3.1.1 The Initial Screen

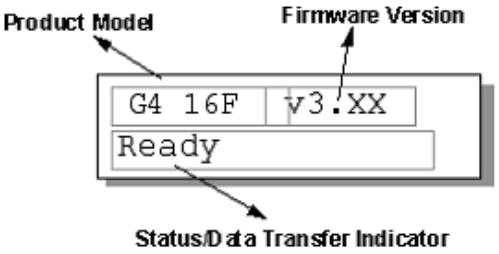

Status/Data Transfer Indicator:

Ready	There is at least one logical drive or logical volume mapped to a host ID/LUN combination.
No Host LUN	No logical drive created or the logical drive has not yet been mapped to any host ID/LUN.
	Indicates the percentage of internal processing resources being consumed, not the host bus throughput. Each block indicates megabytes of data that is currently being processed.

3.1.2 About Logical Drives

In a large enclosure with many drive bays or a configuration that spans across multiple enclosures, including all disk drives into a logical drive MAY NOT BE a good idea. A logical drive with too many members may cause difficulties with maintenance, e.g., rebuilding a failed drive will take a long time.

RAID arrays deliver a high I/O rate by having all disk drives spinning and returning I/O requests simultaneously. If the combined performance of a large array exceeds the maximum transfer rate of a host channel, you will not be able to enjoy the performance gain by simultaneous disk access.

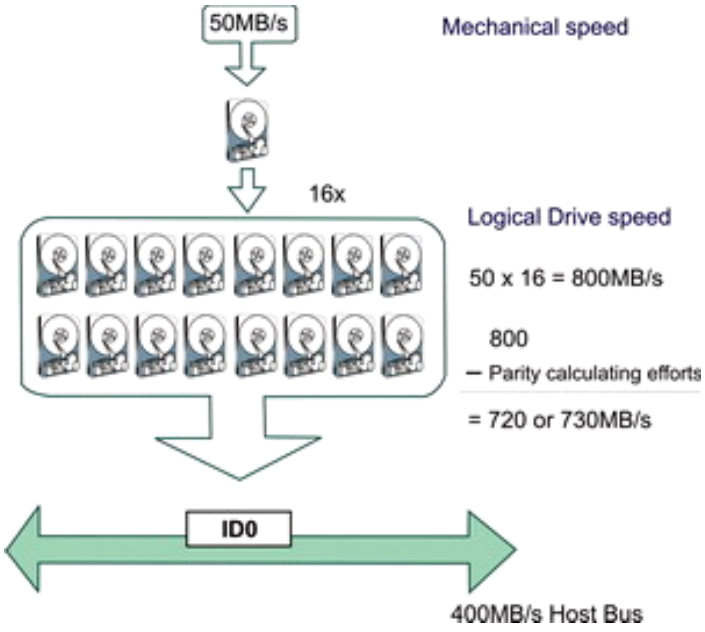

The diagram shows a logical drive consisting of 16 members and associated with a host ID in a 16-bay enclosure. Although host applications may not always generate the theoretical numbers shown here, the host bus bandwidth apparently becomes a bottleneck, and the benefit of simultaneous disk access will be compromised.

3.1.3 Logical Drive Status

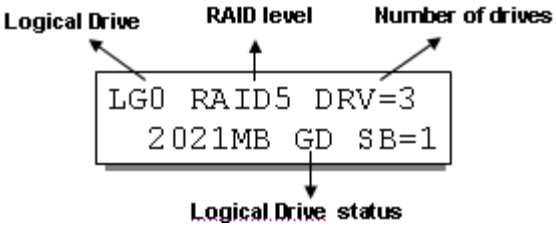

LG#:	The Logical Drive index number.
RAID#:	The RAID level applied for this logical drive.
DRV:	The number of physical drives included in this configuration.
XxxxMB	The capacity of this logical drive.
SB=x	Standby drives available for this logical drive (including Local, Global, and Enclosure Spares). Except the Local spares specifically assigned to other logical configurations, all available spare drive(s) will be counted in this field, including Global and Enclosure-specific Spares.
xxxxMB INITING	The logical drive is now initializing.

xxxxMB INVALID	Fatal failure or incomplete array means that the LD has lost the protection by RAID configuration. If system cannot find some member disks for a specific LD at boot time, the LD will be considered as incomplete. If some member disks of a specific LD fail during operation, the LD will be considered as fatally failed.
xxxxMB GD SB=x	The logical drive is in good condition.
xxxxMB FL SB=x	One member drive failed in this logical drive.
xxxxMB RB SB=x	Logical drive is rebuilding.
xxxxMB DRVMISS	One of the member drives is missing.
INCOMPLETE ARRAY	One or more drives failed in this logical drive.
FATAL FAIL	Two or more member drives failed at the same time, the array is inaccessible
DRV MISS	A member drive is missing, could result from insecure installation
OFF LINE	A logical drive has fatally failed or manually shutdown. This state can result from other faults such as CRC error checksum

3.1.4 Logical Volume Status

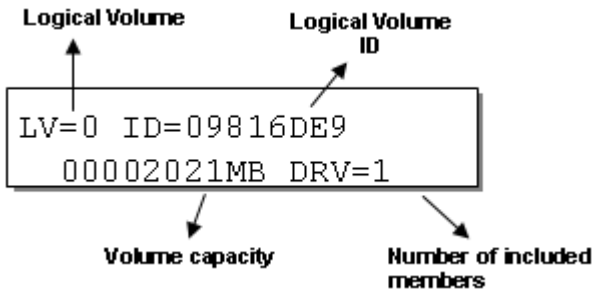

Logical Volume:	The Logical Volume number.
DRV=x:	The number of logical drive(s) included in this logical volume.

Logical Volume ID:	This unique ID is randomly generated by the firmware. In RitePath applications, this ID can be used to identify a RAID volume accessed through two separate host links. Logical drives also have a similar unique ID for ease of identification across a storage network.
xxxMB	The capacity of this logical volume.

3.1.5 Physical Drive Status

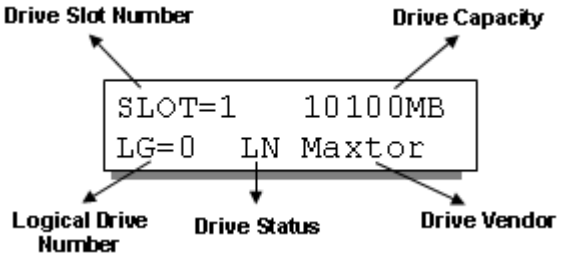

SLOT	The location of this disk drive
LG=*	This drive is a member of logical drive *
LG=x IN	Initializing
LG=x LN	On-line (already a member of a logical configuration)
LG=x RB	Rebuilding
LG=x SB	Local Spare drive
ABSENT	The disk drive does not exist
ADDING	The drive is about to be included in a logical drive through the ADD-Drive procedure
CEDING	When migrating from RAID6 to RAID5, the drive is about to be dismissed from a logical drive. When migration is done, a disbanded drive's status will be indicated as a formatted drive
COPYING	The drive is copying data from a member drive it is about to replace
CLONE	The drive is a clone drive holding the replica of data from a source drive

CLONING	The drive is cloning data from a source drive
EXILED	The drive is considered unreliable, banished from a logical drive, and powered down.

3.1.6 Channel Status

Host Channel

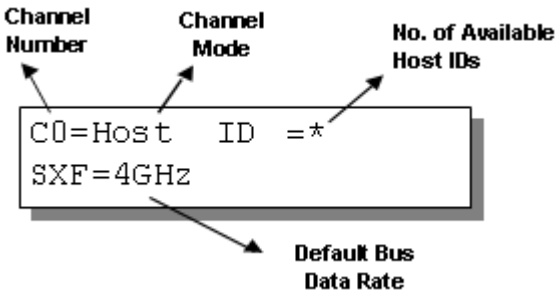

Drive Channel

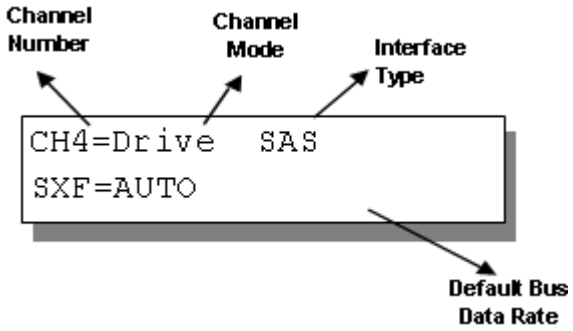

Host	Host channel mode
Drive	Drive channel mode
RCC	Dedicated inter-controller communication channel
AUTO	The default setting is set to the auto-negotiate mode
1 / 1.5 / 2 / 3 / 4 / 6 / 8 / 10 Gbps	Manually configured channel speed
*	Multiple IDs on the channel (Host channel mode only)
(ID number)	IDs are defined as AIDs (Slot A controller IDs) or BIDs (Slot B controller IDs). Slot A is the default location for the Primary RAID controller.

	Host Channel: AIDs or BIDs facilitate the distribution of system workload between RAID controllers that reside in enclosure Slot A and Slot B. An AID and a BID can be associated with the same RAID volume. Drive Channel: A drive channel within a dual-controller configuration will carry both an AID and a BID that are preserved for the channel chip processors on Slot A and Slot B controllers.
NA	No ID applied

NOTE

For a single controller configuration, no IDs will be shown for a drive channel status screen. For a dual-controller configuration, drive channels come with preset IDs. These IDs are assigned to the chip processors on the partner controllers.

3.1.7 Viewing Controller Voltage and Temperature

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Peripheral Dev," then press ENT

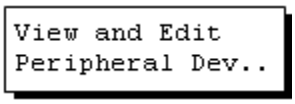

2. Press the up or down arrow keys to select "Ctlr Peripheral Device Config.." Press ENT, choose “View Ctlr Periph Device Status..”, then press ENT.

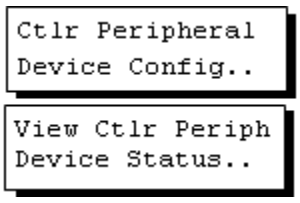

3. Press the up or down arrow keys to choose either “Voltage Monitor” or “Temperature Monitor.”

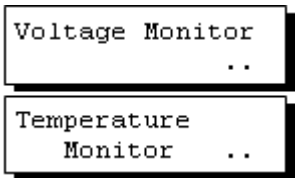

4. Select “Temperature and Voltage Monitor” by pressing ENT.

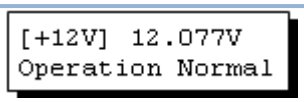

5. Press the up or down arrow keys to browse through the various voltage and temperature statuses.

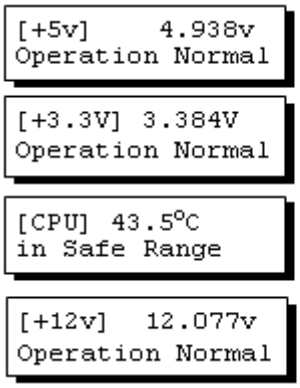

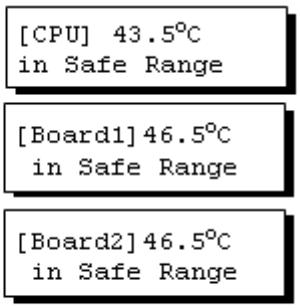

3.1.8 Viewing and Editing Event Logs

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Event Logs," then press ENT.

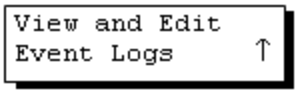

2. Press the up or down arrow keys to browse through the existing event log items.

To see more details about a specific event, use your arrow keys to move to an event, press ENT for 2 seconds to display the first page of event details, then use the arrow keys to move to the next page. When finished reading an event, press the ESC key to return to the event index.

For the limited space on the LCD screen, details of a system event will be displayed in several pages.

To delete a specified item and all events prior to this event, press the ENT key lightly to display the “delete event” confirm message, and then press ENT for 2 seconds to clear the events.

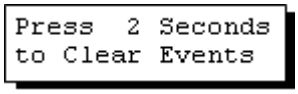

NOTE

The event log will be cleared after the controller is powered off or reset.

Events will be written to the drive reserved space and resetting the subsystem will not erase

the previous event messages.

3.2 Terminal Screen Messages

3.2.1 The Initial Screen

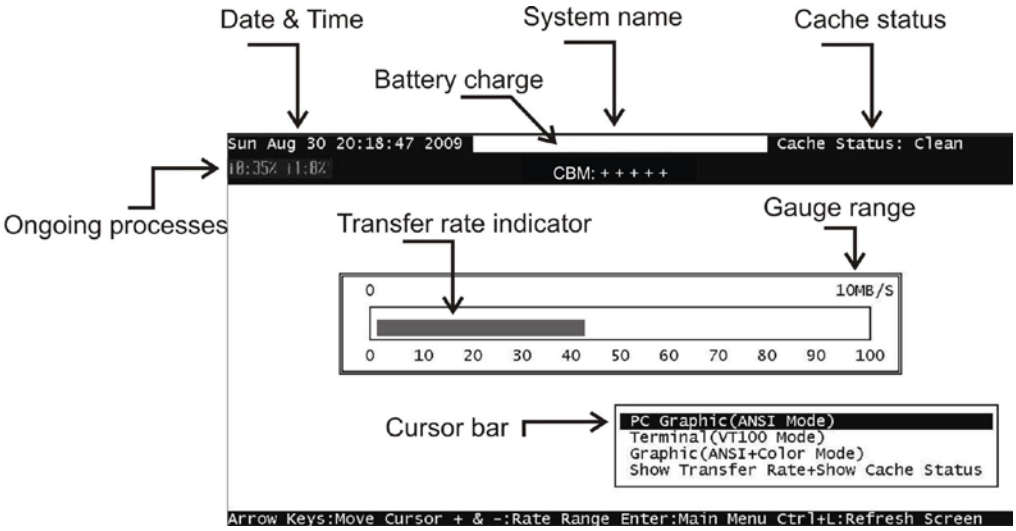

Cursor Bar:□	Highlights the current selection. Move the cursor bar to a desired item, then press [ENTER] to select
Subsystem Name:□	Identifies the type of controller/subsystem or a preset name
Transfer Rate Indicator:□	Indicates the current data transfer rate
Gauge Range:□	Move your cursor bar to “Show Transfer Rate+Show Cache Status.” Press [ENTER] on it to activate the control options, and then use the “Shift” and ”+” or “-“ key combinations to change the gauge range in order to view the transfer rate indicator. The I/O transfer rate will be indicated in percentage against the gauge range. 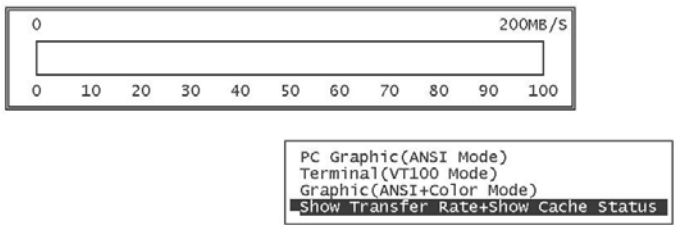 0 200MB/S 0 10 20 30 40 50 60 70 80 90 100 PC Graphic(ANSI Mode) Terminal(VT100 Mode) Graphic(ANSI+Color Mode) Show Transfer Rate+Show Cache Status
Cache Status:□	Indicates current cache status

Write Policy:□	Indicates current write-caching policy
Date & Time:□	Current system date and time, generated by controller real-time clock
PC Graphic (ANSI Mode):□	Enters the Main Menu and operates in ANSI mode
Terminal (VT-100 Mode):□	Enters the Main Menu and operates in VT-100 mode
PC Graphic (ANSI+Color Mode):□	Enters the Main Menu and operates in ANSI color mode
Show Transfer Rate+Show Cache Status:□	Press [ENTER] on this item to show the cache status and transfer rate
Ongoing Processes:	e#: logical drive # is being expanded i#: logical drive # is being initialized R#: logical drive # is being rebuilt P#: logical drive # Parity Regeneration completion ratio S#: logical drive # Media Scan completion ratio For more details, please refer to the Logical Drive Status section in the following discussion.

3.2.2 Main Menu

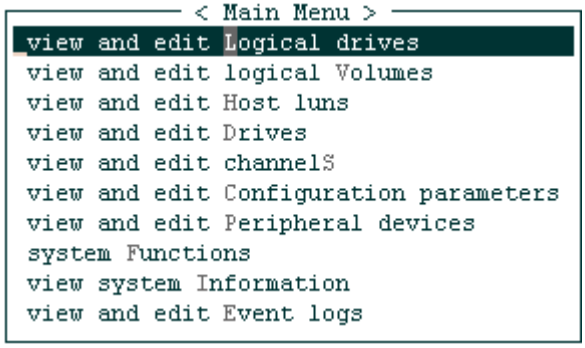
< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channelS
view and edit Configuration parameters
view and edit Peripheral devices
system Functions
view system Information
view and edit Event logs

Use the arrow keys to move the cursor bar through the menu items, then press [ENTER] to choose a menu, or [ESC] to return to the previous menu/screen.

In a subsystem or controller head where battery status can be detected, battery status (CBM on HDX4 models) will be displayed at the top center. Status will be stated as Good, Bad, several “+” (plus) signs (VT-100 mode), or color blocks (ANSI mode) will be used to indicate battery charge. A battery fully-charged status will be indicated by four plus signs (++++) or color blocks.

When initializing or scanning an array, the controller displays progress percentage on the upper left corner of the configuration screen. An “i” indicates array initialization. An “s” stands for scanning process. The number(s) next to them indicate the logical drive index number (e.g., logical drive 0).

3.2.3 Notes on Logical Drives

In a large enclosure with many drive bays or a configuration that spans across multiple enclosures, including all disk drives into a logical drive may not be a good idea. A logical drive with too many members may cause difficulties with maintenance, e.g., rebuild will take a longer time.

RAID arrays deliver a high I/O rate by having all disk drives spinning and returning I/O requests simultaneously. If the combined performance of a large array exceeds the maximum transfer rate of a host channel, you will not be able to enjoy the performance gain by simultaneous disk access.

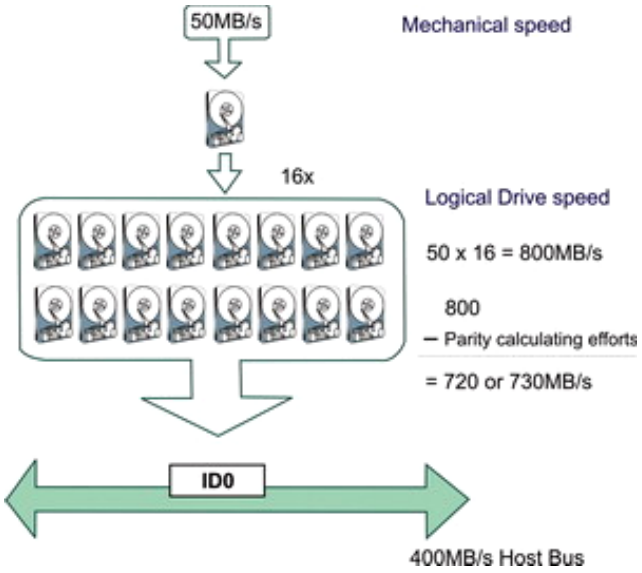

The diagram below shows a logical drive consisting of 16 members and associated with a host ID in a 16-bay enclosure. Although host applications may not always realize the theoretical numbers shown here, the host bus bandwidth apparently becomes a bottleneck, and the benefit of simultaneous disk access will be seriously reduced.

3.2.4 Viewing Logical Drive Status

Go to: View and Edit Logical Drives

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	0	C	#LN	#SB	#FL	NAME
A0	2B699693	0	RAID1	238214	GOOD			7	B	2	1	0	
1													
2													
3													
4													
5													
6													
7													

NOTE

A logical drive in a single-controller subsystem is always managed by one controller, and the “A” or “B” indicator will not appear.

LG	Logical Drive number **A**: Managed by Slot A controller **B**: Managed by Slot B controller
LV	The Logical Volume to which this logical drive belongs
ID	Firmware-generated unique array ID
RAID	RAID level
SIZE (MB)	Capacity of the Logical Drive
Status 1	Logical Drive Status – Column 1 **GOOD**: The logical drive is in good condition **DRV FAILED**: A drive member failed in the logical drive **DRV INITING**: Logical drive is being initialized **INCOMPLETE**: One of the causes of the Incomplete state can be one or more member drives are missing or failed in

	the logical drive **INVALID**: The logical drive was created but has not been fully initialized when another version of firmware is being loaded. After the subsystem resets, the array status should return to normal. Fatal failure or incomplete array means that the LD has lost the protection by RAID configuration. If system cannot find some member disks for a specific LD at boot time, the LD will be considered as incomplete. If some member disks of a specific LD fail during operation, the LD will be considered as fatally failed. **FATAL FAIL**: Two or more member drives failed at the same time, the array is inaccessible **DRV MISS**: A member drive is missing; could result from insecure installation **REBUILDING**: The logical drive is being rebuilt **OFF LINE**: A logical drive has fatally failed or manually shut down. This state can result from other faults such as CRC error checksum
Status 2	Logical Drive Status – Column 2 I: Initializing drives A: Adding drive(s) E: Expanding logical drive H: Add drive operation on hold
Status 3	Logical Drive Status – Column 3 R: Rebuilding the logical drive P: Regenerating array parity
Column O	Logical Drive Status – Stripe size N/A: Default 4: 16KB 5: 32KB

	6: 64KB 7: 128KB 8: 256KB 9: 512KB A: 1024KB
Column C	Logical Drive Status – Write Policy setting B: Write-back T: Write-through
#LN	Total number of drive members in the logical drive
#SB	Standby drives available for the logical drive. This includes all the spare drives (local spare, global spare) available for the specific logical drive
#FL	Number of Failed member(s) in the logical drive
Name	Logical drive name (user configurable)

3.2.5 Viewing Logical Volume Status

Go to: View and Edit Logical Volumes

LV	ID	RAID	Size(MB)	Status 1	2	3	C	#LN	#FL	Name
A0	06FBCA387F22D6EF	RAID0	238214	GOOD			B	1	0	LV0
1										
2										
3										
4										
5										
6										
7										

NOTE

A logical volume in a single-controller subsystem is always managed by one controller, and the “A” or “B” indicator will not appear.

LV	Logical Volume number.
ID	Logical Volume ID number (randomly generated by firmware)
RAID	RAID0 means the members of the logical volume are striped together.
Size(MB)	Capacity of the Logical Volume
#LN	The number of Logical Drive(s) included in this Logical Volume
#FL	The number of failed member(s) within the logical volume.

* For other statuses, please refer to the logical drive information on the previous page.

3.2.6 Viewing Physical Drive Status

Go to: View and Edit Drives

SATA Drives

Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
2	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
3	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
4	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
5	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
6	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
7	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
8	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380

SAS Drives

JBOD	Slot	ChNo	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	1	2	0	238214	300MB	0	ON-LINE	HDT722525DLA380
	2	2	1	238214	300MB	0	ON-LINE	HDT722525DLA380
	3	2	2	238214	300MB	GLOBAL	STAND-BY	HDT722525DLA380
	4	2	3		NONE		ABSENT	
	5	2	4		NONE		ABSENT	
	6	2	5		NONE		ABSENT	
	7	2	6		NONE		ABSENT	
	8	2	7		NONE		ABSENT	

Slot	Drive slot in which a disk drive resides
Size (MB)	Drive capacity
Drive capacity	XXMB Maximum transfer rate of the drive channel interface
ChNo	Channel number. For drives within a SAS expansion enclosure, ChNo will be displayed as n1<n2>, showing two SAS domains connecting to a dual-ported SAS drive.
LG_DRV	The disk drive is a member of logical drive “X.” If the Status column shows “STAND-BY”, the drive is a Local Spare belonging to logical drive “X.”
ID	A logical device ID assigned to the SAS drive.
Status	See the next table.
Vendor and Product ID	The vendor and product model information of the disk drive
JBOD	For disk drives in the expansion enclosures, the number shown in the “JBOD” column indicates which enclosure the disk drives come from. The JBOD ID is configured via DIP switches or a rotary ID switch on the enclosure’s chassis ear.

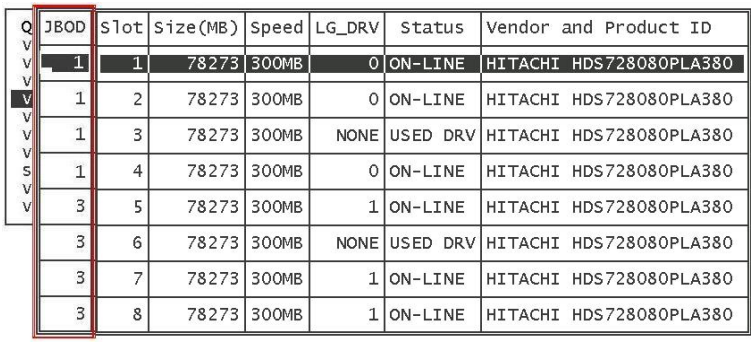

JBOD	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	1	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	2	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	3	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
1	4	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
3	5	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	6	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
3	7	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	8	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380

Status

Global	The disk drive is a Global Spare Drive
INITING	Proceeding with array initialization
ON-LINE	The drive is in good condition
REBUILD	Proceeding with array Rebuild process
STAND-BY	Local Spare Drive or Global Spare Drive. The Local Spare Drive's LG_DRV column will show the logical drive number. The Global Spare Drive's LG_DRV column will show "Global".
NEW DRV	A new drive has not been included in any logical drive or configured as a spare drive
USED DRV	An used drive that is not a member of any logical drive or configured as a spare
FRMT DRV	Formatted drive (drive formatted with a reserved section)
BAD	Failed drive
ABSENT	The disk drive does not exist
ADDING	The drive is about to be included in a logical drive through the ADD-Drive procedure
CEDING	When migrating from RAID6 to RAID5, a member drive is dismissed from the logical configuration. When dismissed from a RAID6 array, the drive status will be indicated as a formatted drive
COPYING	The drive is copying data from a member drive it is about to

	replace
CLONE	The drive is a clone drive holding the replica of data from a source drive
CLONING	The drive is cloning data from a source drive
MISSING	Drive missing (a member drive was once here). This status is shown after boot-up and before I/Os are distributed to the hard drive or accessed by firmware. A missing drive may be corrected by re-inserting the improperly-installed drive tray, etc. If I/Os are distributed and this drive fails to respond, the status will become “failed.”
SB-MISS	Spare drive missing
EXILED	See the description in the next section.

3.2.7 About EXILED Drives

An exiled drive is one that is considered as unreliable by firmware, banished from a logical drive, and then powered down. Banishing and powering down an unreliable drive helps ensure array performance. If a drive is manually disbanded from a logical drive, its status will be indicated as “Exiled.”

An exiled drive state can result from the following:

- Bad drive: A drive fails and is banished from a logical drive.
- Ex-member: If you manually remove and insert a member drive into array, it will be considered as an Exiled drive. Unlike previous firmware, a drive rejoined this way will not become a “Used Drive,” and an automatic rebuild will not start. Note that if a new drive is inserted, rebuild will begin automatically.
- Not Ready:
  A drive that could not be scanned in during the boot process.
  A drive inserted after power-on could not be scanned in.

Other scenarios for change of drive statuses:

- An Exiled drive can be forcefully brought online by removing its 256MB reserved space. Its status will be indicated as “NEW.” However, this method is only recommended for debug purposes.
- An Exiled drive moved to another RAID enclosure will be indicated as a “Used

Drive" because there is no logical drive relationship with it on that enclosure.

- If, for some reasons, a Bad drive can be scanned in after a controller reset, its status will be "Exiled drive" rather than "Used drive."

Same as dealing with a Bad drive, once a drive member turns into an Exiled drive, firmware automatically rebuilds a logical drive if a hot-spare is available. If a hot-spare is not available, you should replace the Exiled drive as soon as possible.

3.2.8 Viewing Channel Status

Go to: View and Edit Channels

Fibre-to-SATA Configuration

Chl	Mode	ID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	112	Auto	Serial	F		4.0 GHz	Serial
1	Host	112	Auto	Serial	F		4.0 GHz	Serial
2	Drive		Auto	SATA			3.0 GHz	SATA
3	Drive		Auto	SATA			3.0 GHz	SATA
4	Drive		Auto	SATA			3.0 GHz	SATA
5	Drive		Auto	SATA			3.0 GHz	SATA
6	Drive		Auto	SATA			3.0 GHz	SATA
7	Drive		Auto	SATA			3.0 GHz	SATA

Fibre-to-SAS Configuration

Chl	Mode	AID	BID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	112	NA	Auto	Serial	F		4.0 GHz	Serial
1	Host	NA	113	Auto	Serial	F		4.0 GHz	Serial
2	Drive			3.0 GHz	SAS	A		3.0 GHz	SAS
3(4)	Drive			3.0 GHz	SAS	A		3.0 GHz	SAS
4(3)	Drive			3.0 GHz	SAS	A		3.0 GHz	SAS
5(C)	RCCOM			3.0 GHz	SATA			3.0 GHz	SATA
6(C)	RCCOM			3.0 GHz	SATA			3.0 GHz	SATA
7(C)	RCCOM			3.0 GHz	SATA			3.0 GHz	SATA

Chl — Channel number; expansion links are also defined as drive channels yet with a bracketed number showing the counterpart SAS domain (in a dual-controller configuration).

Mode	Channel mode RCCOM: Redundant controller communication channel Host: Host Channel mode Drive: Drive Channel mode
AID	IDs managed by the Slot A controller *: Multiple IDs were applied (Host Channel mode only) (ID number): Host Channel: Specific IDs managed by the Slot A controller for host LUN mapping Drive Channel: Specific ID reserved for the channel processor on the Slot A controller
BID	IDs managed by the Slot B controller *: Multiple IDs were applied (Host Channel mode only) (ID number) Host Channel: Specific IDs managed by the Slot B controller for host LUN mapping Drive Channel: Specific ID reserved for the channel processor on the Slot B controller; used in redundant controller mode NA: No channel ID applied
AUTO	Channel bus data rate set to auto speed negotiation
DefSynClk	Default bus synchronous clock: ??.?GHz The default setting of the channel is ??.?GHz in Synchronous mode. Async.: The default setting of the channel is Asynchronous mode.
DefWid	Default bus width: Serial: Serial transfer protocol; for Fibre Channel or SAS Channel
S	Signal:

	F: Fibre A: SAS
Term	Terminator Status: (not applied here in Fibre-to-SAS/SATA solutions) On: Terminator is enabled. Off: Terminator is disabled. Diff: The channel is a Differential channel. The terminator can only be installed/removed physically. Empty: Non-SCSI bus
CurSynClk	Current bus synchronous clock: ??.?GHz: The default setting of the channel bus is ??.? GHz Async.: The default setting of the channel bus is Asynchronous mode. (empty): The default bus synchronous clock has changed. Reset the controller for the changes to take effect.
CurWid	Current Bus Width: Serial: Serial transfer protocol; Fibre Channel, SAS Channel, SATA Channel.

3.2.9 Viewing Controller Voltage and Temperature

Go to: View and Edit Peripheral Devices > Controller Peripheral Device Configuration > Voltage and Temperature Parameters

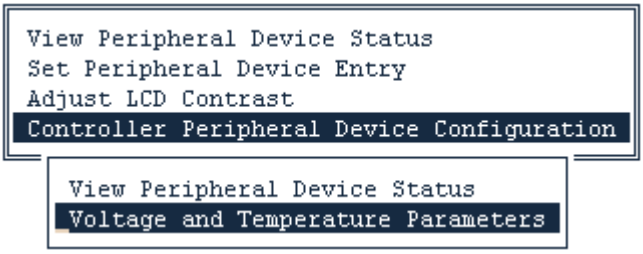

ITEM	VALUE	STATUS
+3.3V Value	3.368 V	Operation Normal
+5V Value	5.18 V	Operation Normal
+12V Value	12.32 V	Operation Normal
CPU Temperature	55.0 (C)	Within Safe Range
Controller Temperature(1)	42.0 (C)	Within Safe Range
Controller Temperature(2)	51.0 (C)	Within Safe Range

The current status of voltage and temperature detected by the controller will be displayed on-screen and will be stated as normal, out of order, or within the safety range.

3.2.10 Viewing Event Logs on Screen

Go to: View and Edit Event Logs

When errors occur, you may want to trace the records to see what has happened to your system. The controller's event log management records all events starting from the time when the system is powered on, recording up to 1,000 events. Powering off or resetting the controller will automatically delete all of the recorded event logs.

The event logs are stored in disk reserved space, and hence the event logs are available after system reset. Disk reserved space is automatically created when composing a logical drive. With no logical drives, event logs can not be preserved.

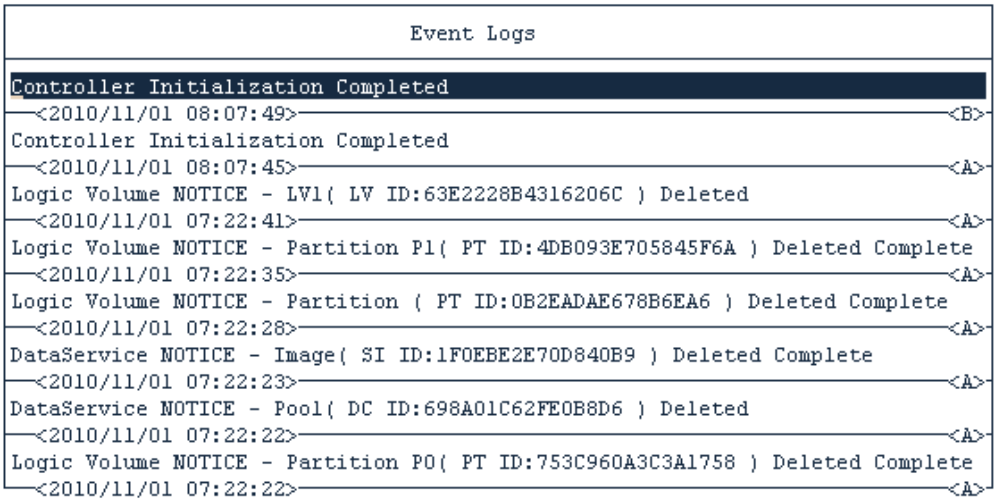

To check for more details about a specific event, move the cursor bar to highlight a specific event and press the [Space] key to display the complete event information.

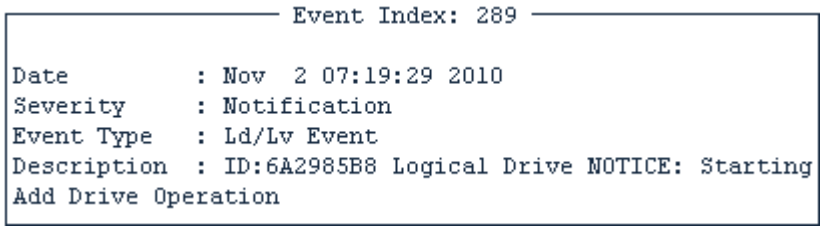
Event Index: 289
Date : Nov 2 07:19:29 2010
Severity : Notification
Event Type : Ld/Lv Event
Description : ID:6A2985B8 Logical Drive NOTICE: Starting
Add Drive Operation

To clear the saved event logs, scroll the cursor down to select an event and press [ENTER] to delete the event and the events below.

Choose Yes to clear the recorded event logs.

4 Optimizing & Preparing Tasks

There are preference parameters that cannot be easily altered after the creation of logical arrays. Reconfiguration takes time and inappropriate configurations prevent you from getting the best performance from your Galaxy arrays. It is therefore highly recommended to thoroughly consider preferences such as stripe sizes, caching parameters, etc. before creating your logical arrays.

4.1 Configuring Caching Parameters

4.1.1 Deciding the Stripe Size

Each RAID level has a preset value for the array stripe size. If you prefer a different stripe size for a RAID array (a logical drive), you must backup or move the stored data elsewhere and re-create the array.

Listed below are the default stripe sizes implemented with different RAID levels. These values should be adequate for optimal performance with most applications.

Level	Stripe Size:
RAID0	128KB
RAID1	128KB
RAID3	16KB
RAID5	128KB
RAID6	128KB
NRAID	128KB

Stripe sizes different from the above defaults can be manually applied to individual logical drives during the initial configuration stage to match the access sizes conducted by your host applications.

NOTE

The Stripe size here refers to the “Inner Stripe Size” specifying the chunk size allocated on each individual data drive for parallel access instead of the “Outer Stripe

Size" which is the sum of chunks on all data drives.

According to Berkeley paper, strip size is a single chunk of data written into a member disk drive. The Stripe size option here is the strip size.

Although stripe size can be adjusted on a per logical drive basis, users are not encouraged to make a change to the default values.

Smaller stripe sizes are ideal for I/Os that are transaction-based and randomly accessed. However, using the wrong stripe size can cause problems. For example, when an array set at 16KB stripe size receives files of 128KB size, each drive will have to spin and write many more times to conduct small fragment 16KB writes to hard disks.

4.1.2 Enabling Write-Back Cache

Go to: View and Edit Configuration Parameters > Caching Parameters > Write-Back Cache

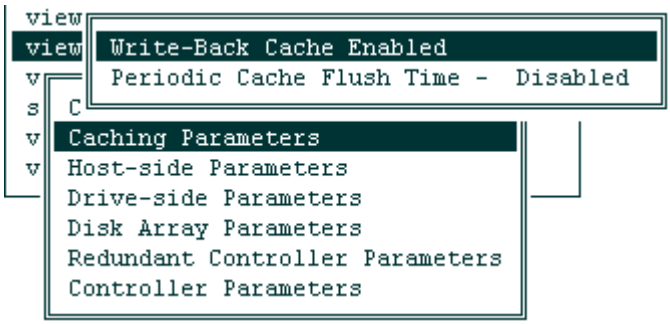

As one of the sub-menus in "Caching Parameters," this option controls the cached write policy.

When "Write-back" is "Enabled," the write requests from the host will be held in cache memory and distributed to disk drives later. When "Write-back" is "Disabled" (i.e., the Write-through is adopted,) host writes will be directly distributed to individual disk drives. Select Yes in the dialog box that follows to confirm the setting.

- The Write-through mode is safer if your controller is not configured in a redundant pair and there is no battery backup or UPS device to protect cached data.
- Write-back caching can dramatically improve write performance by caching the unfinished writes in memory and letting them be committed to drives in a more efficient manner. In the event of power failure, a battery backup module can hold cached data for days. On the HDX4 series, a CBM module will keep cached data

in its flash memory and hence there is no concern with the hold-up time (usually 72 hours) by the batteries.

4.1.3 Enabling Write-Back

The Write-back options can be found either here in the Configuration Parameters menu or in the “View and Edit Logical Volume” sub-menu (logical drive or logical volume). The Write-back option found here is the system general setting that applies to all logical volumes. If you apply a different write-back mode for individual logical volumes, then those logical volumes will operate with its write-back mode regardless of the system’s general setting.

1. From the Main Menu, select “View and Edit Config Parms,” “Caching Parameters,” and press ENT.

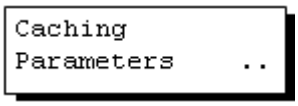

2. As one of the sub-menus in "Caching Parameters," this option controls the cached write function. Press ENT to enable or disable “Write-back Cache.”

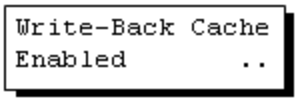

3. Press ENT for two seconds to confirm. The current status will be displayed on the LCD.

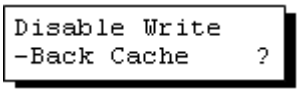

4. The Write caching options also appear in array-specific (logical drive and logical volume) configuration menu and should look like the screens shown below.

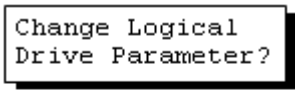

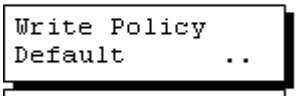

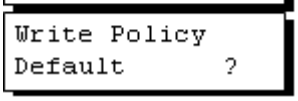

4.1.4 Trigging Events

System General Setting

Go to: View and Edit Configuration Parameters > Caching Parameters > Write-Back Cache

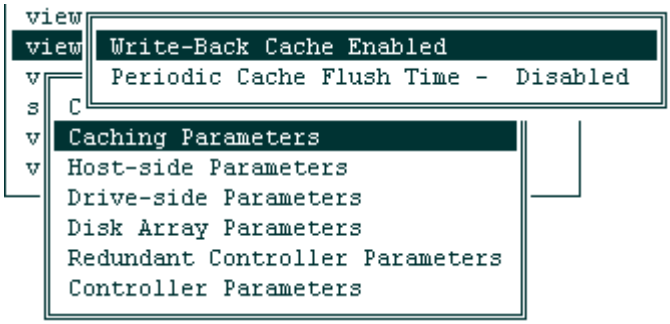

Array-Specific Setting

Go to: View and Edit Logical Volumes > (Logical Volume) > Write Policy

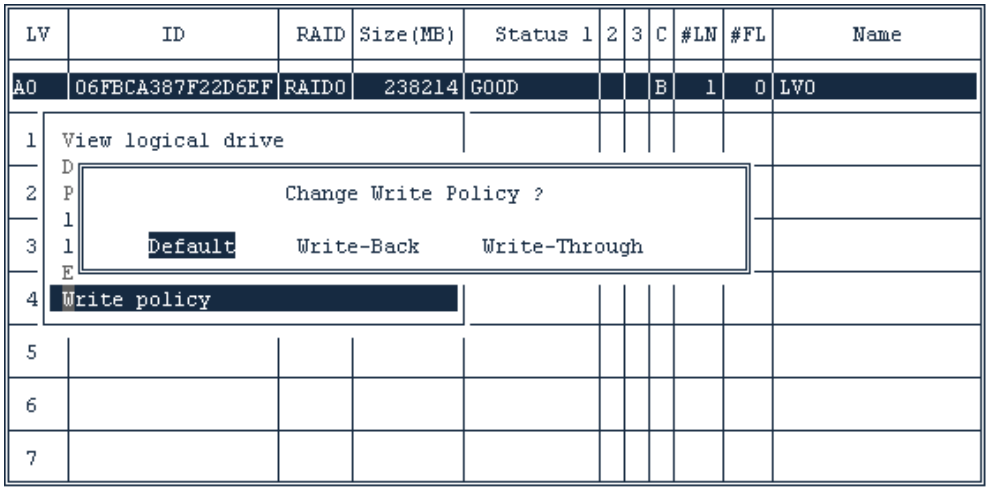

The configuration options are related to the Event Triggered Operation feature.

1. The Event Triggered Operation feature allows the firmware to automatically enable or disable Write-back caching in the event of component failure or critical system alarms.

2. As shown below, a relatively unsafe condition will force the controller to assume a conservative “Write-through” caching mode, e.g., a SPU or cooling fan failure.

```
                          Event Logs
NOTICE: Controller BBU Fully Charged !
<==<Fri Jun 24 02:09:06 2005>==========================================<B>-
NOTICE:Controller Default Write Policy Restore
<==<Thu Jun 23 23:02:32 2005>==========================================<A>-
NOTICE:Force Controller Write-Through on Trigger Cause
<==<Thu Jun 23 23:02:26 2005>==========================================<A>-
NOTICE:Controller BBU is Charging !
<==<Thu Jun 23 23:02:21 2005>==========================================<B>-
NOTICE: Controller BBU Fully Charged !
<==<Tue Jun 21 19:40:45 2005>==========================================<A>-
NOTICE:Controller Default Write Policy Restore
<==<Tue Jun 21 16:38:49 2005>==========================================<A>-
NOTICE:Force Controller Write-Through on Trigger Cause
<==<Tue Jun 21 16:38:42 2005>==========================================<A>-
NOTICE:Controller BBU is Charging !
<==<Tue Jun 21 16:38:39 2005>==========================================<A>-
```

3. A ”Default“ Write-back option is available with individual logical arrays. Default means the logical drive will follow the system’s general caching configuration.

4. If a logical array’s Write-back mode is set to “Default,” the caching mode of that particular array will be dynamically controlled by the firmware.

5. If the Write-back mode is manually specified as “Enabled” or “Disabled” in a particular logical array, then I/Os directed to that array will be handled in accordance with the setting regardless of the system’s general setting.

Event Trigger configurations

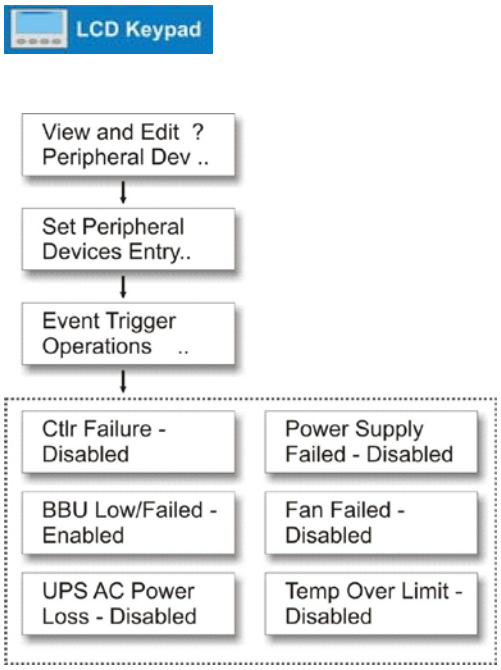
LCD Keypad
View and Edit ? Peripheral Dev ..
Set Peripheral Devices Entry..
Event Trigger Operations ..
Ctlr Failure - Disabled
Power Supply Failed - Disabled
BBU Low/Failed - Enabled
Fan Failed - Disabled
UPS AC Power Loss - Disabled
Temp Over Limit - Disabled

Terminal

Go to: View and Edit Peripheral Devices > Set Peripheral Device Entry > Event Trigger Operations

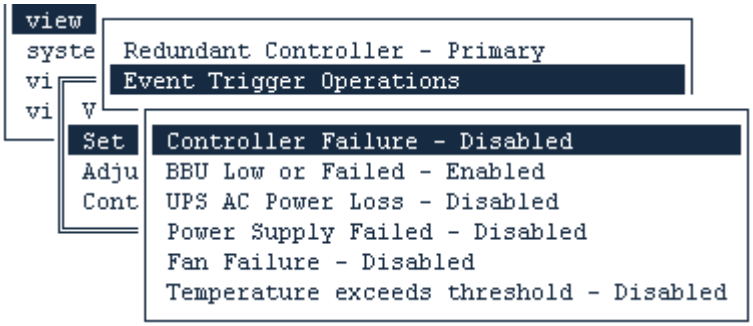

Enable one or more preferred options on the list to protect your array from hardware faults.

4.1.5 Flushing Cache Periodically

If Write-back caching is preferred for better performance yet data integrity is also a concern, e.g., a configuration without battery protection or synchronized cache between partner controllers, the system can be configured to flush the cached writes at preset intervals.

LCD Keypad

1. From the Main Menu, select “View and Edit Config Parms,” “Caching Parameters,” and press ENT.

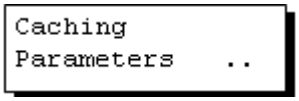

2. Use the arrow keys to scroll through the options and select “Periodic CachFlush Time”, and then press ENT to proceed.

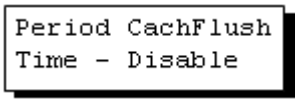

3. The “Set Cache Flush Time – Disable” appears. The default is “Disable.” Use your arrow keys to select an option from “ConSync,” “30sec,” to “600 sec.” “ConSync” stands for “continuously synchronized.”

```
Set Cache Flush
Time - Disable
```

4. Press ENT to select and press ESC to exit and the setting will take effect immediately.

Terminal

Go to: View and Edit Configuration Parameters > Caching Parameters > Periodic Cache Flush Time

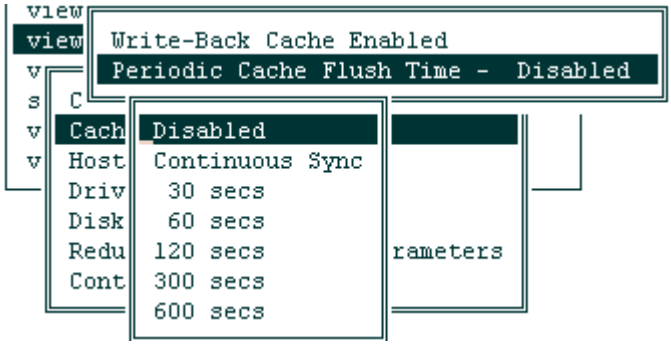

Note that the “Continuous Sync” option holds data in cache for as long as necessary to complete a write operation and immediately commits a write request to hard drives if it is not followed by a series of sequential writes.

NOTE

Every time you change the Caching Parameters you must reset the controller for the changes to take effect.

4.2 Preparing Channels and Channel IDs

4.2.1 Notes on Channel Mode Settings

Go to: View and Edit Channels > (Channel)

The Galaxy HDX4 subsystems come with preset data paths and there is no need to modify channel modes.

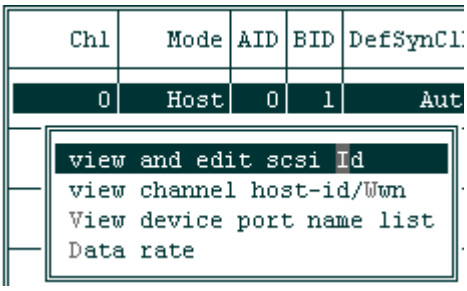

- For different channel assignments, please refer to the Hardware manual that came with your subsystem. For example, an HDX4 FC<>FC system comes with 6 FC channels. 2 of the FC channels can be connected to host or serve as drive loops.
- Technical terms like Slot A, Slot B, RCC (Redundant Controller Communications), and DRVRCC will only appear in a dual-controller configuration.
- The latest Galaxy HDX4 models come with dedicated RCC (Redundant Controller Communications) chipsets that provide communication paths strung between partner RAID controllers. The “Drive+RCC” and “RCC” options will not appear on the list of available channel modes.
- You can still find these RCC channels on the channel list, only that there are no configurable options with these dedicated RCC paths.
- Most Galaxy HDX4 RAID subsystems have preset host or drive channels interfaced through a backplane. The channel mode options are not available on these models.

4.2.2 Configuring Channel ID

Each host channel comes with a default AID (one that managed by controller A) and/or a BID, which will not be sufficient if your subsystem comes in a complex dual-active controller configuration.

In a dual-active controller configuration, you need to manually create more Slot A or Slot B Channel IDs to distribute the workload between partner RAID controllers. The idea is diagrammed below:

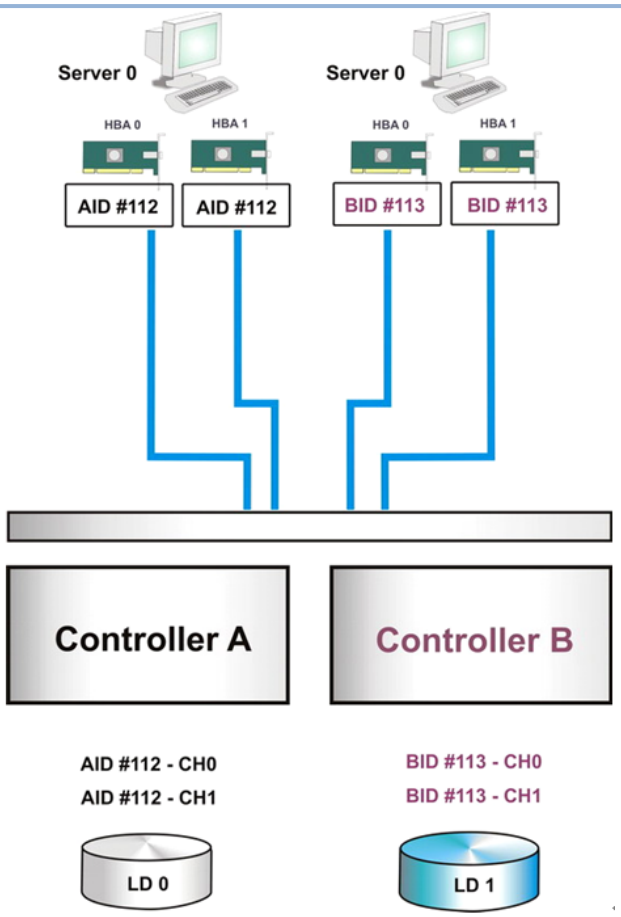

Configuration:

- 2 logical drives (LD)
- LD LUN mapping associations:
- LD0: CH0 AID112 & CH1 AID112
- LD1: CH0 BID113 &CH1 BID113

Controller B IDs need to be manually created. Note that in this example, a multi-pathing software is required to manage the fault-tolerant links to a RAID storage volume.

A logical group of physical drives can be associated either with Controller A IDs or Controller B IDs through the host LUN mapping process. These A or B IDs then appear to the application servers as storage capacity volumes. As a rule of thumb, a logical drive associated with AIDs is managed by Controller A. One that is associated with BIDs is managed by Controller B.

You also need to assign logical volumes to an individual controller. The option is found in View and Edit Logical Volume.

Depending on how many RAID capacity volumes you wish visible to your application servers, create one or more Controller A or Controller B IDs. In firmware menus, these IDs are specified as the slot A or slot B IDs. You may also present storage volumes to host using the LUN numbers under channel IDs. A max. of 1024 LUNs are supported, and up to 32 LUNs under each ID.

In the event of a single controller failure, IDs managed by the failed controllers will be taken over and managed by the surviving controller.

NOTE

The HDX4 supports the cross-controller ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs. However, mapping to both controllers' IDs is only beneficial when it is difficult making fault-tolerant host links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. Currently, SAS switch is not popular on the market. For Fibre-host systems, fault-tolerant links can easily be made with external bypass such as Fibre Channel switches.

For details of fault-tolerant link connections, please refer to your system Hardware Manual.

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Channels," then press ENT.

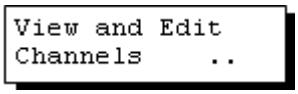

2. Channel information will be displayed. Press ENT on the host channel you wish the ID changed.

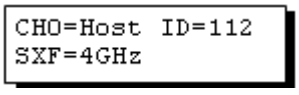

3. Press the up or down arrow keys to select "Set Channel ID," then press ENT.

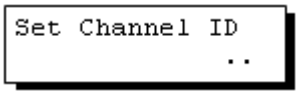

4. Use the up or down arrow keys to browse through the existing host IDs. Press ENT on any ID combination to continue.

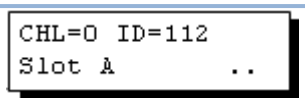

Terminal

Go to: View and Edit Channels > (Channel) > View and Edit SCSI ID

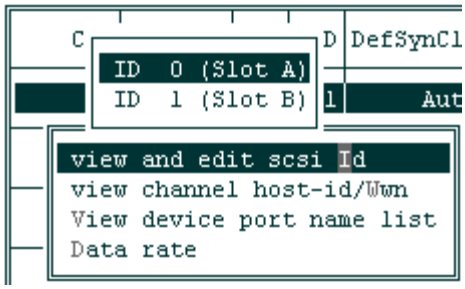

1. Select a host channel, press [ENTER] to display the command list.
2. Select “View and Edit ID.” A list of existing ID(s) will be displayed on the screen. As a default, the subsystem comes with only a Slot A controller ID.
3. Select one of the existing IDs and press [ENTER]. You may then add a new ID or delete an existing ID.

4.2.3 Adding a Host ID

LCD Keypad

1. Press ENT on a host channel, on “Set Channel ID”, and then on an existing ID.
2. Use the up and down arrow keys to select “Set Channel ID", then press ENT.

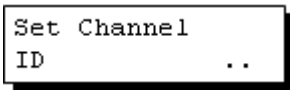

3. An existing ID displays.

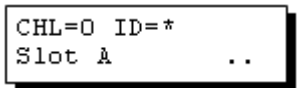

4. Press ENT to display “Add Channel ID.” Press ENT again to display the question mark.

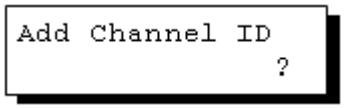

5. In a dual-controller configuration, once you enter the Add ID process, use the up and down arrow keys to select either the Slot A or Slot B controller.

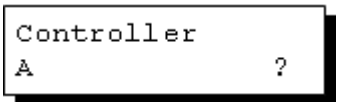

6. An ID next to the existing ID will display on the screen. Use arrow keys to select an ID. When the preferred ID is selected, press ENT for two seconds to complete the process.

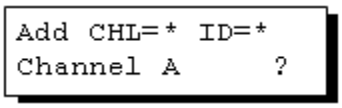

7. A prompt will remind you to reset the subsystem for the configuration change to take effect. You may press ENT to reset the subsystem immediately or you may press ESC to continue adding other host IDs and reset the subsystem later.

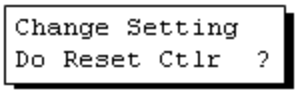

Go to: View and Edit Channels > (Channel) > View and Edit SCSI ID > (ID) > Add Channel > Slot > (ID)

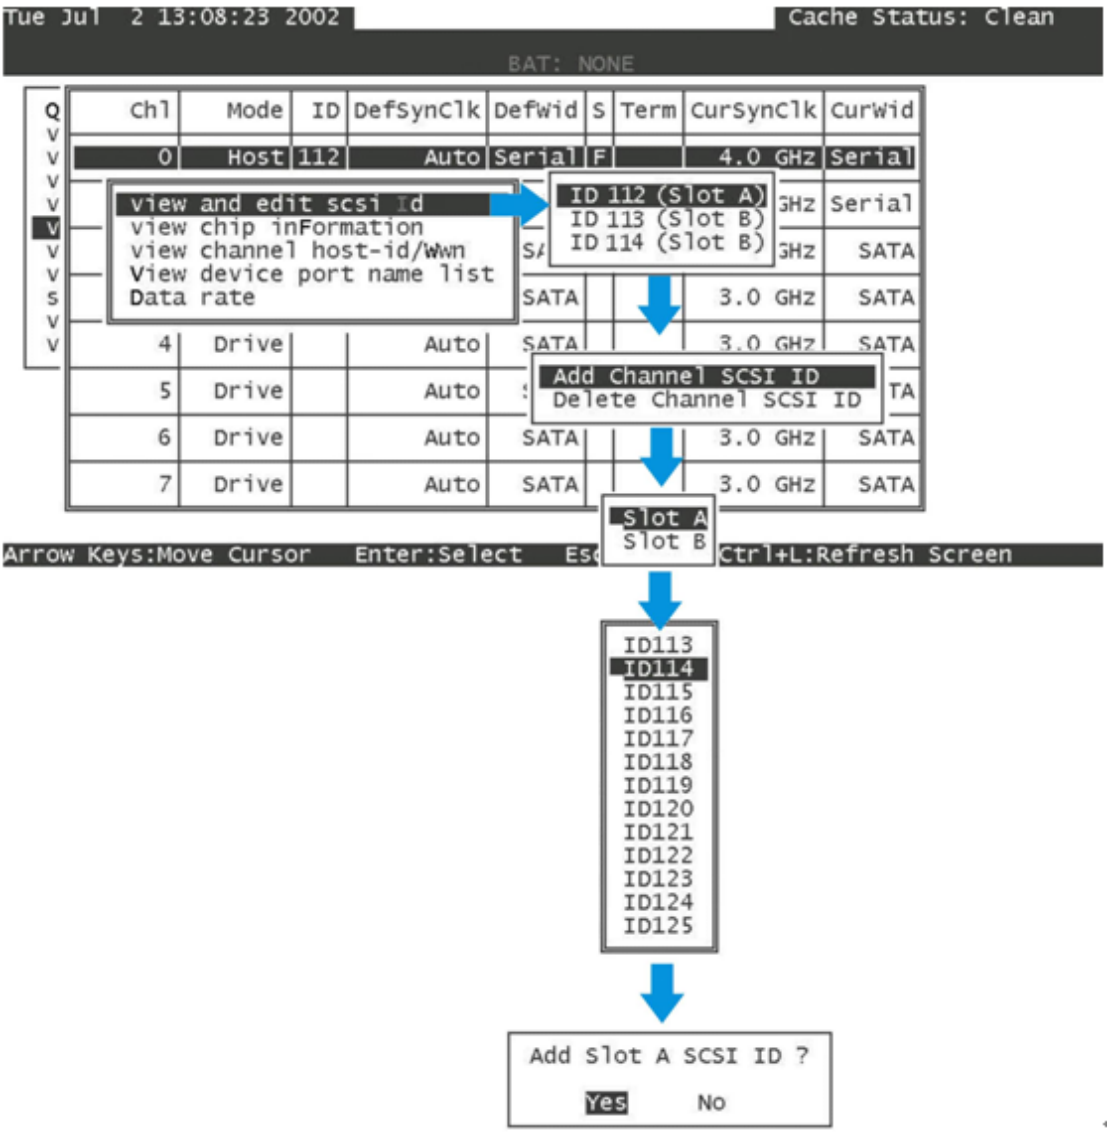

1. Press [ENTER] on one of the existing IDs.
2. Select Add Channel ID.
3. Specify the host ID either as the Slot A or Slot B ID. Press [ENTER] to proceed.
4. Available IDs will appear in a pull-down list. Select by pressing [ENTER] and then select Yes to confirm.
5. A confirmation box will prompt to remind you to reset the controller for the configuration to take effect. You may select Yes for an immediate reset or No to reset later.

```
NOTICE: Change made to this setting will NOT take effect
        until the controller is RESET. Prior to resetting the
        controller, operation may not proceed normally.

          Do you want to reset the controller now ?

                    Yes                    No
```

4.2.4 Deleting an ID

LCD Keypad

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Channels," then press ENT.

```
View and Edit
Channels      ↕
```

2. The first host channel should appear. Press ENT to select a host channel.

```
CHL=Host ID=113
SXF=AUTO    ..
```

3. Press ENT on “Set Channel ID..”

```
Set Channel
ID          ..
```

4. A list of host channel and host ID combinations will appear. Use the up or down arrow keys to select the ID you wish to remove. Press ENT to select a channel ID combination.

```
CHL=0 ID=112
              ..
```

5. You will then be prompted by the “Add Channel ID” option. Press the down arrow key to proceed.

```
Add Channel ID
              ..
```

6. The “Delete Channel ID” option will appear. Press ENT to display the confirmation box. Press ENT for two seconds to remove the ID.

```
Delete Channel
ID          ..
```

7. A prompt will remind you to reset the subsystem for the configuration change to take effect. You may press ENT to reset the subsystem immediately or you may press ESC to continue adding other host IDs and reset the subsystem later.

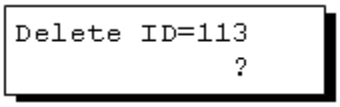

Terminal

Go to: View and Edit Channels > (Channel) > View and Edit SCSI ID > (ID) > Delete Channel

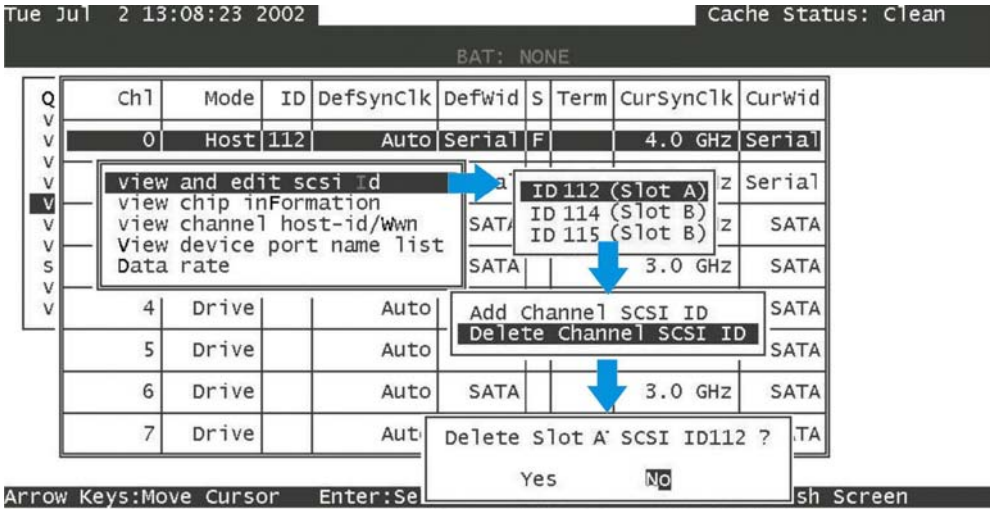
Tue Jul 2 13:08:23 2002
Cache Status: Clean
BAT: NONE
Chl Mode ID DefSynClk DefWid S Term CurSynClk CurWid
0 Host 112 Auto Serial F 4.0 GHz Serial
view and edit scsi Id
view chip inFormation
view channel host-id/Wwn
View device port name list
Data rate
ID 112 (Slot A)
ID 114 (Slot B)
ID 115 (Slot B)
Add Channel SCSI ID
Delete Channel SCSI ID
4 Drive Auto SATA
5 Drive Auto SATA
6 Drive Auto SATA 3.0 GHz SATA
7 Drive
Delete Slot A SCSI ID112 ?
Yes No
Arrow Keys:Move Cursor Enter:Se
sh Screen

NOTE

Every time you change a channel ID, you must reset the subsystem/controller for the changes to take effect.

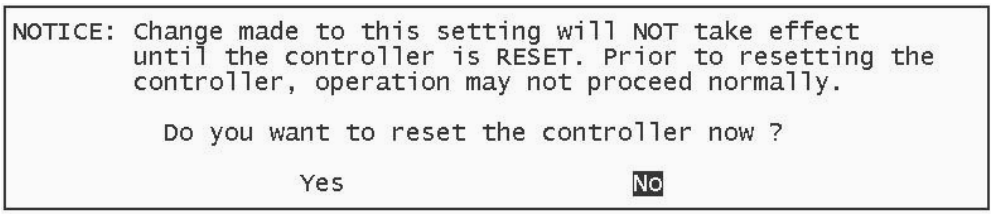
NOTICE: Change made to this setting will NOT take effect
until the controller is RESET. Prior to resetting the
controller, operation may not proceed normally.
Do you want to reset the controller now ?
Yes
No

At least one controller's ID should be present on each channel bus.

4.2.5 Selecting the Data Rate (Host Channel Bus)

The data rate default is "AUTO" and should work fine with most configurations. In some cases, you may want to install a 8Gbps interface subsystem in a storage

network consisting of 4Gbps devices. Please note that mixing 8G and 4G devices in a storage network may not be supported with all kinds of HBAs or Fibre switches.

Note that the data rate setting on drive channel menus is the maximum transfer rate of the channel bus in that mode. It does not mean a single disk drive can actually carry out that amount of sustained read/write performance.

LCD Keypad

1. From Main Menu, select “View and Edit Channels,” and then the host channel you wish to change.

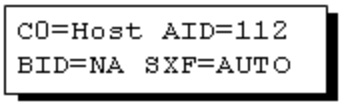

2. Press ENT on the channel and use the arrow keys to find the “Data Rate” option.

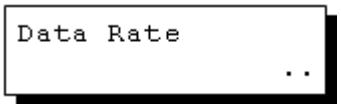

3. Press ENT on the Data Rate option to display “Set Chl=X Data Rate To AUTO?”, where “X” stands for the channel number.

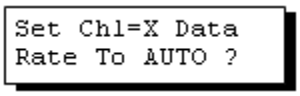

4. Use your arrow keys to display the desired data rate. Press ENT to confirm the selection.

4.2.6 Selecting the Data Rate (Drive Channel)

LCD Keypad

1. From Main Menu, select “View and Edit Channels,” and then the drive channel you wish to change.

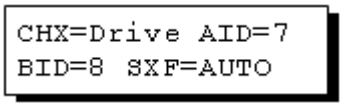

2. Press ENT on the channel and use the arrow keys to find the “Data Rate” option.

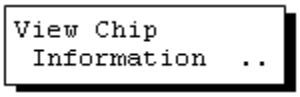

3. Press ENT on the Data Rate option to display “Set Chl=X Data Rate To AUTO?”, where “X” stands for the channel number.

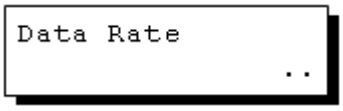

4. Use your arrow keys to display a data rate value which ranges from 33 to 300MBps (SATA drive channels). Press ENT to confirm a selection.

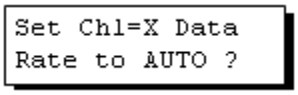

Go to: View and Edit Channels > (Drive Channel) > Data Rate

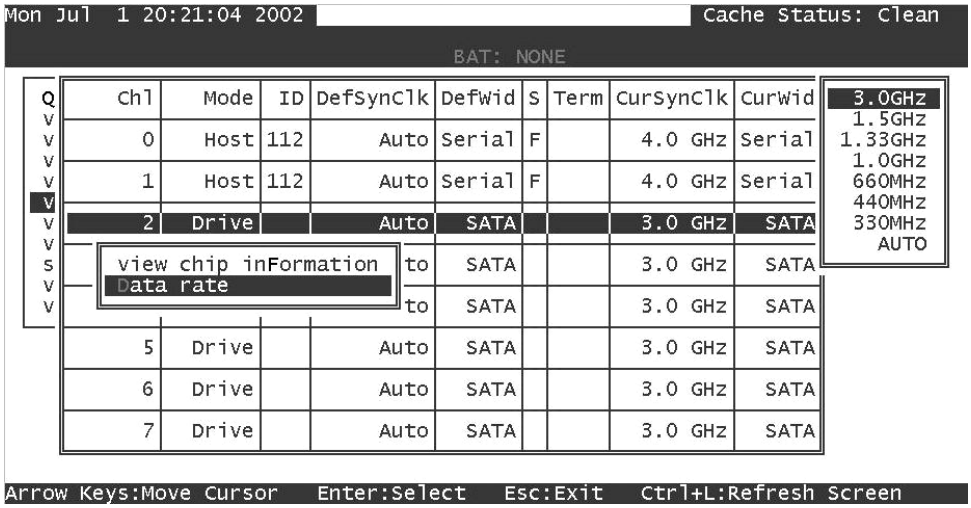
Mon Jul 1 20:21:04 2002
Cache Status: Clean
BAT: NONE
Chl Mode ID DefSynClk DefWid S Term CurSynClk CurWid
0 Host 112 Auto Serial F 4.0 GHz Serial
1 Host 112 Auto Serial F 4.0 GHz Serial
2 Drive Auto SATA 3.0 GHz SATA
view chip information
Data rate
5 Drive Auto SATA 3.0 GHz SATA
6 Drive Auto SATA 3.0 GHz SATA
7 Drive Auto SATA 3.0 GHz SATA
3.0GHz
1.5GHz
1.33GHz
1.0GHz
660MHz
440MHz
330MHz
AUTO
Arrow Keys:Move Cursor Enter:Select Esc:Exit Ctrl+L:Refresh Screen

Note that the Galaxy HDX4 series does not support SATA disk drives at a 1.5Gb/s speed. Some SATA drives may come with a default set to 1.5Gb/s. Use a drive’s jumpers or configuration utility to change its setting.

4.3 Setting Controller Date and Time

Setting the correct date and time is important especially when tracing system faults or applying automated maintenance utilities such as Media Scan scheduler.

Galaxy’s latest Galaxy Array Manager software supports time synchronization with SNTP time server and it is recommended to specify your time zone.

4.3.1 Selecting the Time Zone

The controller uses GMT (Greenwich Mean Time), a 24-hour clock. To change the clock to your local time zone, enter the numbers of hours earlier or later than the Greenwich Mean Time after the plus (+) or minus (-) sign. For example, “+9” is Japan’s time zone.

1. Choose “View and Edit Configuration Parameters,” “Controller Parameters," then press ENT.

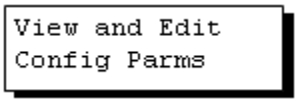

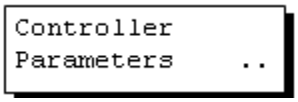

2. Press the up or down arrow keys to scroll down and select “Set Controller Date and Time”, then press ENT.

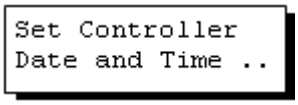

3. Choose “Time Zone” by pressing ENT.

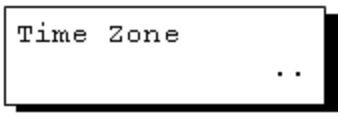

4. Use the down key to enter the plus sign and the up key to enter numbers.

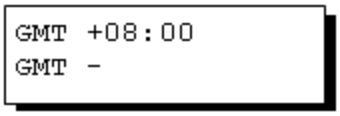

Go to: View and Edit Configuration Parameters > Controller Parameters > Set Controller Date and Time > Time Zone

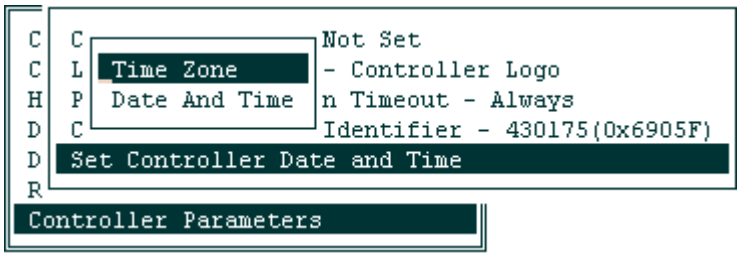

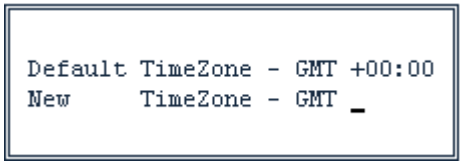

4.3.2 Setting the Date and Time

LCD Keypad

1. Use your arrow keys to scroll down and select “Date and Time” by pressing ENT.

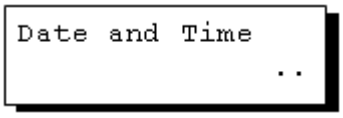

2. Use the arrow keys to select and enter the numeric representatives in the following order: month, day, hour, minute, and the year. Use up/down arrow keys to change the number displayed on screen, press ENT to shift to the next number.

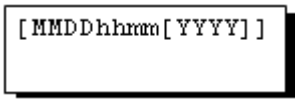

Terminal

Go to: View and Edit Configuration Parameters > Controller Parameters > Set Controller Date and Time > Date and Time

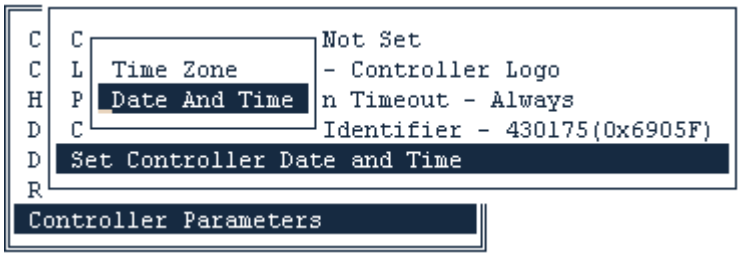

```
Time and Date [MMDDhhmm[YYYY]] : _
```

Enter time and date in its numeric representatives in the following order: month, day, hour, minute, and the year.

4.4 Detecting Faulty Drives

When enabled, the Auto Rebuild check time scans the drive bus/channel on which a failed drive resides. If the drive swap check detects a replacement drive, the system firmware will automatically proceed with the array rebuild process.

Without the Auto Rebuild check time, the rebuild process can be manually initiated through a “rebuild” command under the “View and Edit Logical Drive” sub-menu. This check time mechanism is specifically applicable in a configuration where no hot-spare is available.

1. Select “View and Edit Config Parms” from the terminal Main Menu. Enter its sub-menus by pressing ENT.

```
View and Edit
Channels      ↕
```

2. Use arrow keys to select “Drive-side Parameters.” press ENT to enter its sub-menus.

```
Drive-side
Parameters    ..
```

3. There are a dozen configurable options under Drive-side parameters. Use arrow keys to select “Auto Rebuild on Drv Swap.” Press ENT on it to change the setting. The options range from Disabled and 5 to 60 seconds.

```
Auto Rebuild on
Drv Swap-15 sec
```

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Auto Rebuild on Drive Swap

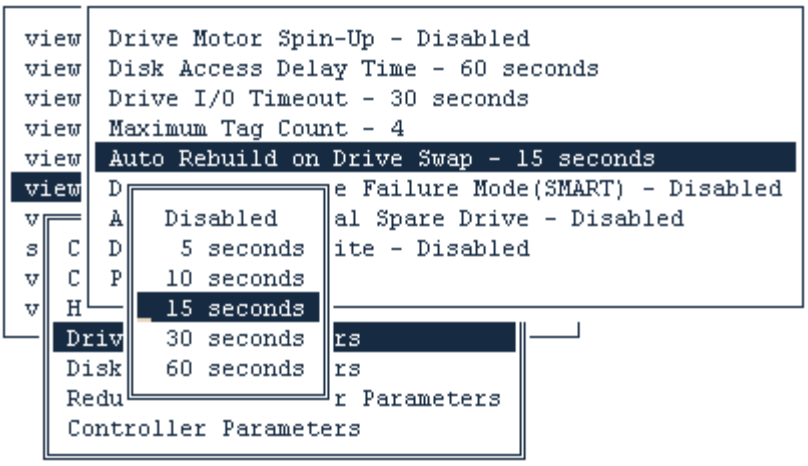

4.5 Assigning Spare Drives

Shown below are two Spare drive policies designed to prevent configuration errors: Auto-assign Global Spare and Enclosure Spare Drive.

NOTE

The capacity of spare drives must be equal to or greater than that of member drives.

4.5.1 About Auto-assignment of a Global Spare

The Auto-Assign Global Spare feature is designed to reduce the chance of down time by operator's negligence. Shown on the left is a RAID enclosure with its drives configured into two arrays and a Global Spare. One logical drive consists of 8 members; the other consists of 7.

Diagrams below show how the Auto-assign mechanism helps prevent downtime:

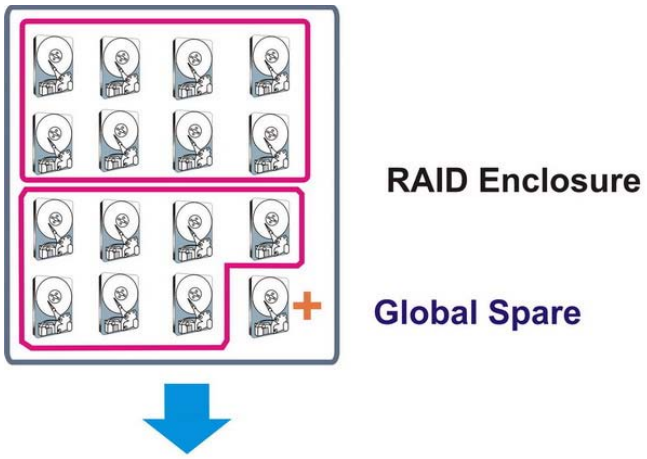

1. A member drive in one of the two logical drives fails. The Global Spare immediately participates in the rebuild.

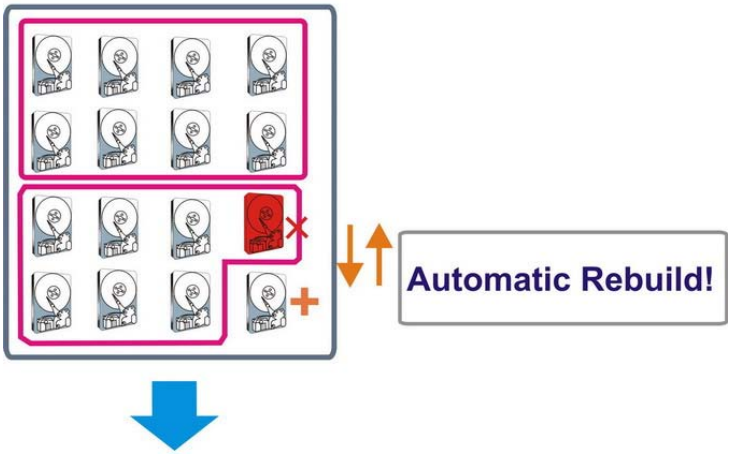

2. The Failed drive is then replaced by a replacement drive. The original Global Spare becomes a member of the 7-drive array.

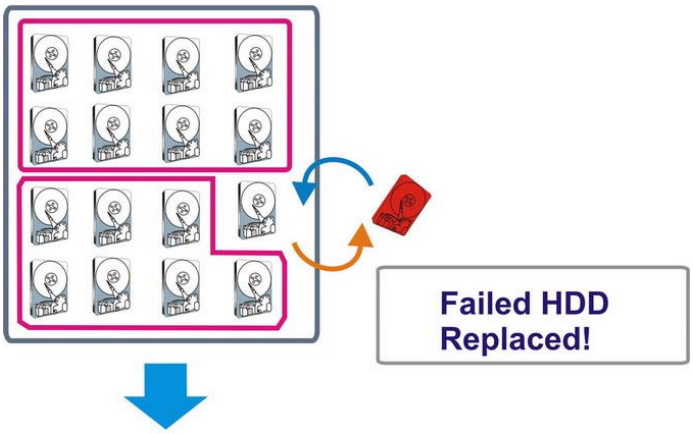

3. With the Auto-Assign feature, firmware automatically configures the replacement drive as a Global Spare. The Auto-Assign feature prevents the situation when a failed drive is replaced and the system administrator forgets to configure the

replacement drive as another Global Spare leaving the array vulnerable to the occurrence of another drive failure.

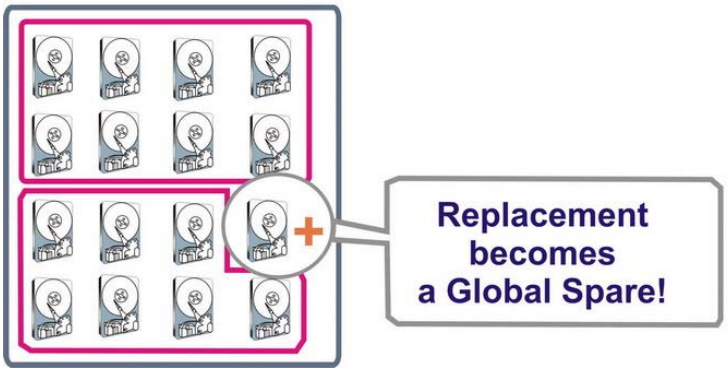

4.5.2 Auto-Assigning a Global Spare

LCD Keypad

1. Select “View and Edit Config Parms” from the terminal Main Menu. Enter its sub-menus by pressing ENT.

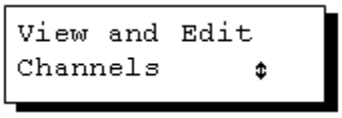

2. Use arrow keys to select “Drive-side Parameters.” press ENT to enter its sub-menus.

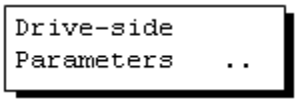

3. There are a dozen of configurable options. Use the arrow keys to select “Periodic SAF-TE ChkTime -.” Press ENT on it to change the setting. The options ranges from Disabled, 50ms,… to 60 seconds.

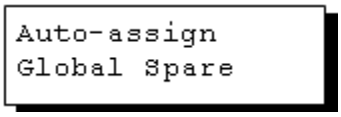

Terminal

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Auto-Assign Global Spare Drive

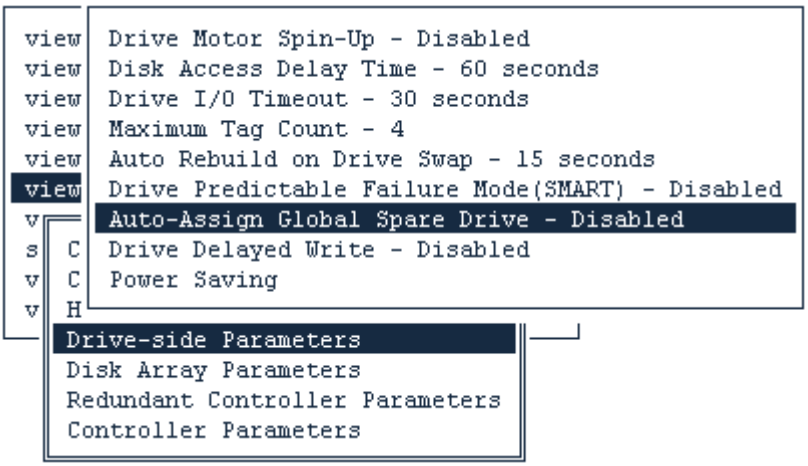

4.5.3 About Enclosure Spare

In addition to the traditional “Local” and “Global” hot spares, another hot-spare type “Enclosure” is added. Global hot-spare may cause a problem as diagrammed below in a storage application consisting of multiple enclosures:

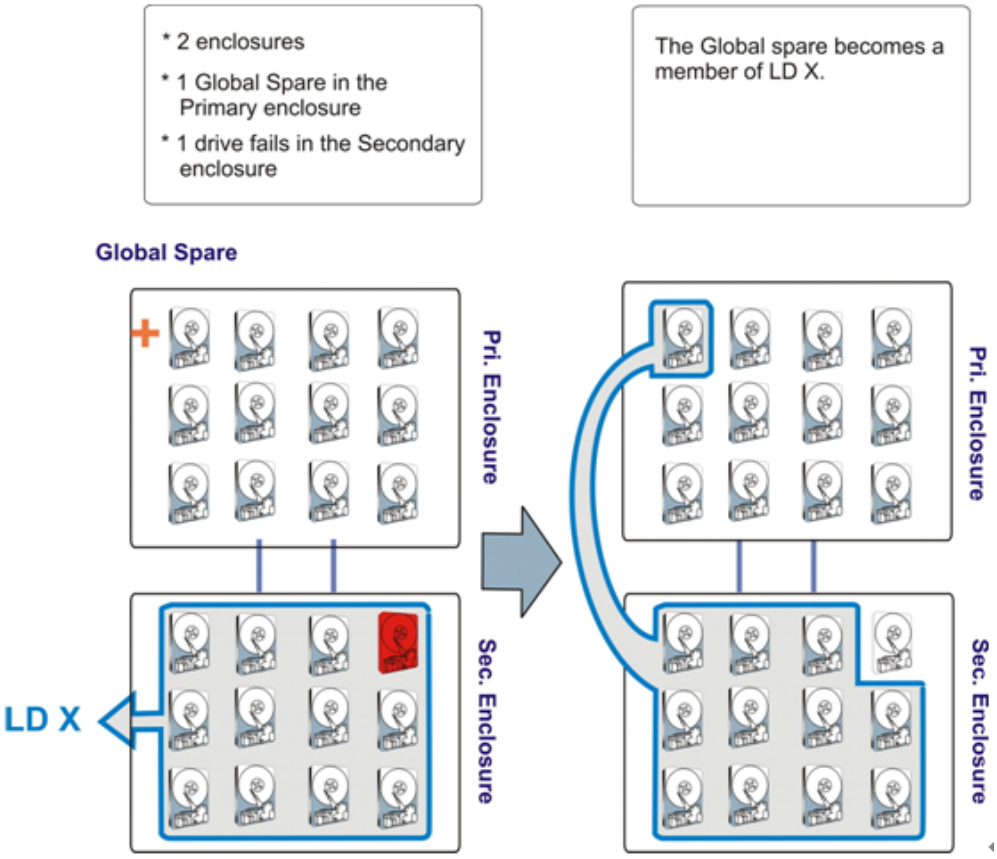
* 2 enclosures
* 1 Global Spare in the Primary enclosure
* 1 drive fails in the Secondary enclosure
The Global spare becomes a member of LD X.
Global Spare
Pri. Enclosure
Sec. Enclosure
LD X
Pri. Enclosure
Sec. Enclosure

A Global spare participates in the rebuild of any failed drive. When a Global spare participates in the rebuild of a logical drive in another enclosure, it will become the member of that logical drive. Although the logical drive can work properly, however, spanning a logical configuration across different enclosures increases the chance of

removing the wrong drive, accidentally mixing SAS and SATA drives of different RPM's, etc.

The Enclosure Spare helps prevent the situation from causing inconvenience. An Enclosure Spare only participates in the rebuild of drives that reside in the same enclosure.

4.5.4 Assigning an Enclosure Spare

1. Select "View and Edit Drives" from the terminal Main Menu. Enter its sub-menus by pressing ENT.

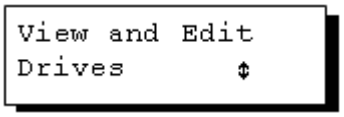

2. Use arrow keys to select a new or formatted drive. Press ENT on it to display drive-specific functions.

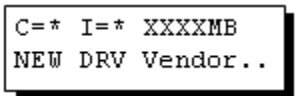

3. Use arrow keys to find "Add Enclosure Spare Drive." Press ENT on it for two seconds to confirm.

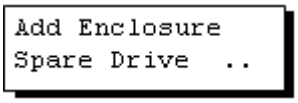

4. A message prompts to confirm a successful configuration. Press ESC to skip the message

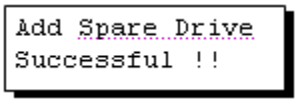

5. The disk drive should now be indicated as an Enclosure spare.

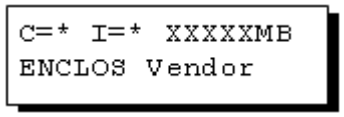

Go to: View and Edit Drives > (Non-assigned drive) > Add Enclosure Spare Drive

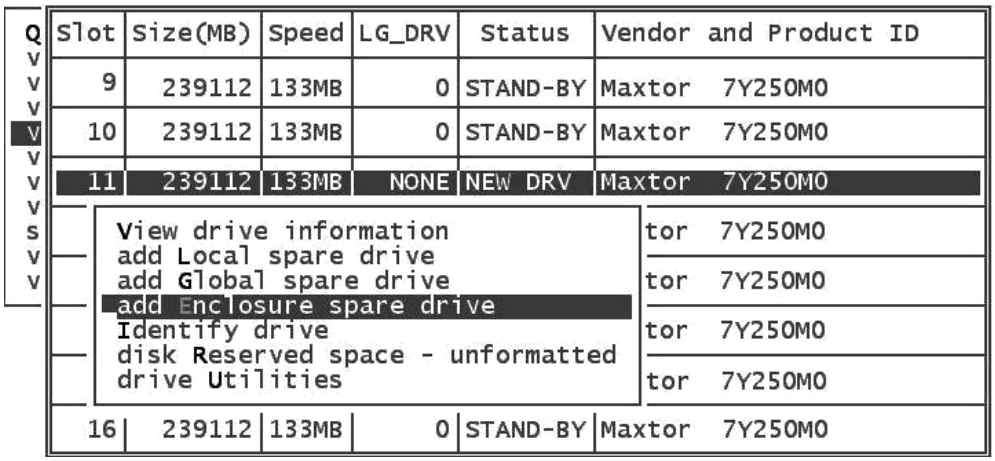

4.6 Enabling Delayed Write to Drive

This option applies to disk drives that come with embedded read-ahead or writer buffers. When enabled, the embedded buffer can improve read/write performance. However, this option should be disabled for mission-critical applications. In the event of power outage or drive failures, data cached in drive buffers may be lost, and data inconsistency will occur. For performance-oriented applications, this option can be enabled.

Following are the defaults for different storage configurations:

- On dual-controller models that come with BBUs, the default is “Disabled.”
- On single-controller models that come without BBUs, the default is “Enabled.”

LCD Keypad

6. Select “View and Edit Config Parms” from the terminal Main Menu. Enter its sub-menus by pressing ENT.

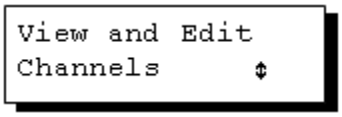

7. Use arrow keys to select “Drive-side Parameters.” press ENT to enter its sub-menus.

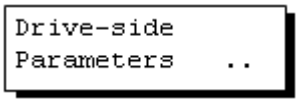

8. There are a dozen of configurable options. Use the arrow keys to select “Drive Delayed Write -.” Press ENT on it to change the setting.

```
Drive Delayed
Write - Disabled
```

Terminal

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Drive Delayed Write

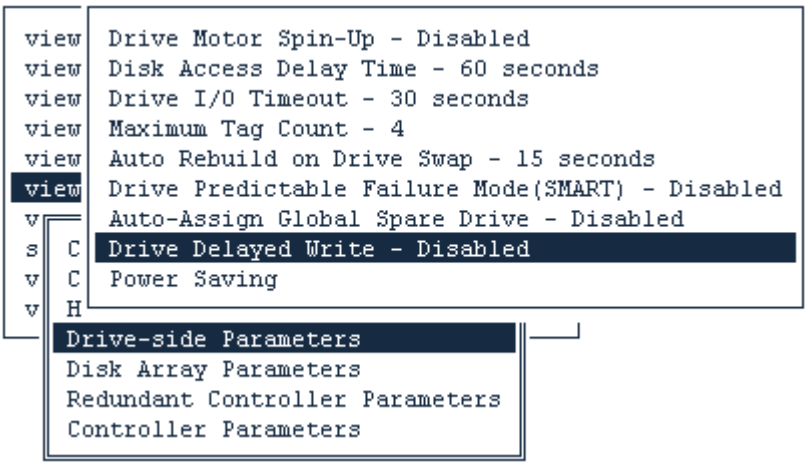
view Drive Motor Spin-Up - Disabled
view Disk Access Delay Time - 60 seconds
view Drive I/O Timeout - 30 seconds
view Maximum Tag Count - 4
view Auto Rebuild on Drive Swap - 15 seconds
view Drive Predictable Failure Mode(SMART) - Disabled
Auto-Assign Global Spare Drive - Disabled
Drive Delayed Write - Disabled
Power Saving
Drive-side Parameters
Disk Array Parameters
Redundant Controller Parameters
Controller Parameters

4.7 Configuring System Functions

Choose “System Functions” in the Main Menu, then press ENT. Press the up or down arrow keys to select a submenu, then press ENT.

4.7.1 Muting Beeper Sound

LCD Keypad

When the controller’s beeper has been activated, choose “Mute Beeper," then press ENT to turn the beeper off temporarily for the current event. The beeper will still activate on the next event. A mute button can also be found on the LCD keypad panel.

```
Mute Beeper
                ..
```

Terminal

Go to: System Functions > Mute Beeper

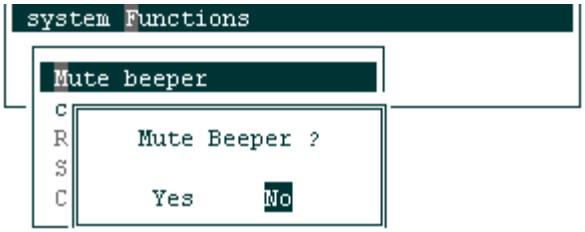

When the subsystem's beeper (onboard alarm) is activated, choose "Mute Beeper," then press [ENTER]. Choose Yes and press [ENTER] in the next dialog box to turn the beeper off temporarily for the current event. The beeper will still be activated by the next event.

4.7.2 Changing the Password

Use the controller's password to protect the system from unauthorized access. Once the controller's password is set, regardless of whether the LCD panel, the RS-232C terminal interface or the Galaxy Array Manager is used, an user can only configure and monitor the RAID controller by providing the correct password.

NOTE

The controller requests a password whenever a user is entering the main menu from the initial screen or a configuration change is made. If the controller is going to be left unattended, the "Password Validation Timeout" should be set to "Always Check." The controller password and controller name share a 32-character space. The maximum number of characters for a controller password is 32. If 31 characters are used for a controller name, there will be only one character left for the controller password and vice versa. The current firmware revisions support a 32-character name space.

LCD Keypad

1. To set or change the controller password, press the up or down arrow keys to select "Change Password," then press ENT.

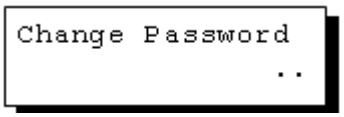

2. If the password has previously been set, the controller will ask for the old password first. If password has not yet been set, the controller will directly ask for the new password. The password cannot be replaced unless the correct old password is provided.

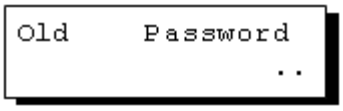

3. Press the up or down arrow keys to select a character, then press ENT to move to the next space.

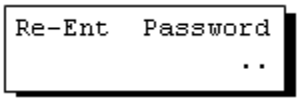

4. After entering all the characters (alphabetic or numeric), press ENT for two seconds to confirm. If the password is correct, or there is no preset password, it will ask for the new password. Enter the password again to confirm.

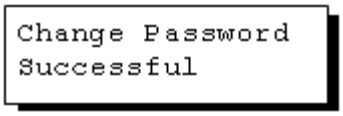

Go to: System Functions > Change Password

To disable or delete the password, press ENT on the first flashing digit for two seconds when requested to enter a new password. The existing password will be deleted. No password checking will occur when entering the Main Menu from the initial terminal screen or making configuration changes.

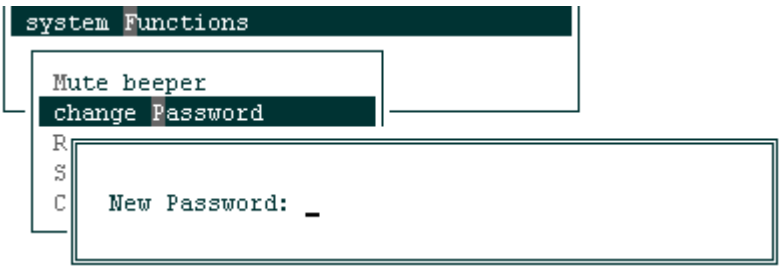

If a password has previously been set, the controller will ask for the old password first.

If the password has not yet been set, the controller will directly ask for the new password. The password cannot be replaced unless the correct old password is provided.

Key-in the old password, then press [ENTER]. If the password is incorrect, it will not allow you to change the password. Instead, it will display the message “Password incorrect!,” then return to the previous menu.

If the password is correct, or there is no preset password, it will request for a new password.

Enter the desired password in the column, then press [ENTER]. The next dialog box will display “Re-Enter Password.” Enter the password again to confirm and press [ENTER].

The new password will now become the controller’s password. Providing the correct password is necessary when entering the Main Menu from the initial screen.

4.7.3 Resetting the Controller

1. To reset the controller without powering off the system, press the up or down arrow keys to “Reset Controller,” then press ENT.

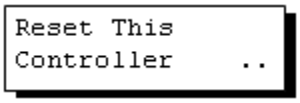

2. Press ENT again for two seconds to confirm. The controller will now reset.

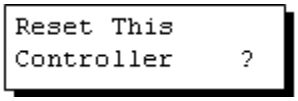

Go to: System Functions > Reset Controller

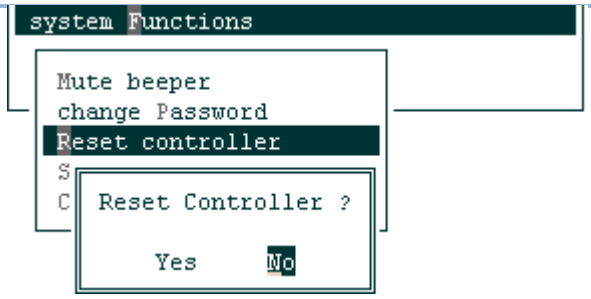

NOTE

Before resetting or powering off the RAID controller (subsystem) it is advised you execute the Shutdown Controller function to flush the cache contents in the memory in order to reduce the chance of encountering data inconsistency.

4.7.4 Shutting Down the Controller

Before powering off the controller, unwritten data may still reside in cache memory. Use the “Shutdown Controller” function to flush the cache content.

NOTE

This function does NOT mean shutting down the server.

LCD Keypad

1. Press the up or down arrow keys to “Shutdown Controller,” then press ENT. Press ENT again for two seconds to confirm.

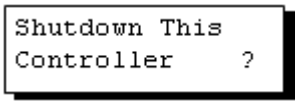

2. The controller will now flush the cache memory. Press ENT for two seconds to confirm and to reset or power off the subsystem.

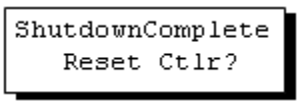

Terminal

Go to: System Functions > Shutdown Controller

Before powering off the controller, unwritten data may still reside in cache memory. Use the “Shutdown Controller” function to flush the cache content.

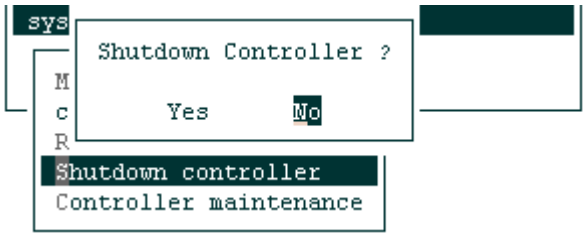

For Controller Maintenance functions, such as “Download Firmware,” please refer to Appendix B.

4.7.5 Saving NVRAM to Disks

You can choose to backup your controller-dependent configuration information to disks. We strongly recommend using this function to save the configuration profile whenever a configuration change is made. The information will be distributed to every logical drive in the RAID system. If using the Galaxy Array Manager, you can save your configuration details as a file to a computer system drive.

Here are some notes for saving NVRAM.

- The "Save NVRAM" function can be used to preserve you system configuration or to duplicate system configurations to multiple storage systems. However, the logical drive mapping will not be duplicated when downloading the NVRAM contents of one system to another. LUN mapping adheres to specific “name tags” of logical drives, and therefore you have to manually repeat the LUN mapping process. All of the download functions will prompt for a file source from the current workstation.
- The Save NVRAM function keeps a record of all configuration data in firmware, including host-side, drive-side, logical drive configurations, and controller-related preferences.
- Data Service settings, e.g., Snapshot configuration, will not be preserved by the Save NVRAM function. The snapshot meta table is kept on the drive media of a source volume.
- A RAID configuration of drives must exist for the controller to write NVRAM content onto it.

3. From the Main Menu, choose “System Functions.” Use arrow keys to scroll down and select “Controller Maintenance,” “Save NVRAM to Disks,” then press ENT.

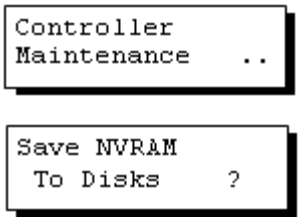

4. Press ENT for two seconds on the message prompt, “Save NVRAM to Disks?”

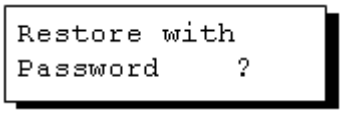

5. A prompt will inform you that NVRAM information has been successfully saved.

Terminal

Go to: System Functions > Controller Maintenance > Export NVRAM to Reserved Space

At least a RAID configuration must exist for the controller to write your configuration data onto it.

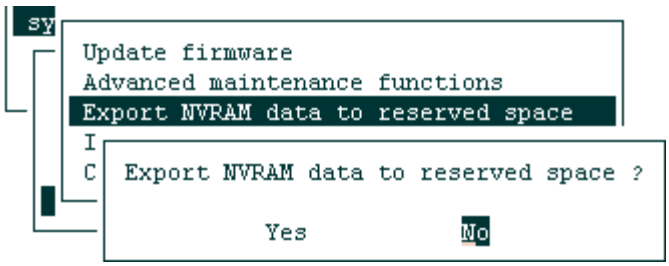

4.7.6 Restoring NVRAM from Disks

If you want to restore your NVRAM information that was previously saved onto the array, use this function to restore the configuration setting.

LCD Keypad

1. From the Main Menu, choose “System Functions.” Use arrow keys to scroll down and select “Controller Maintenance,” “Restore NVRAM from Disks..,” and then press ENT.

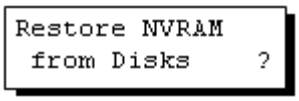

2. Press ENT for two seconds to confirm.

3. In case your previous password (reserved at the time you saved your NVRAM configuration contents) is different from your current password, you are provided with the options whether to restore the password you previously saved with your configuration profile.

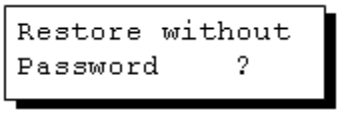

4. A prompt will inform you that the controller NVRAM data has been successfully restored from disks.

Terminal

Go to: System Functions > Controller Maintenance > Import NVRAM Data from Reserved Space

In case your previous password (preserved at the time you saved your NVRAM configuration contents) is different from your current password, you are provided with the options whether to restore the password you previously saved.

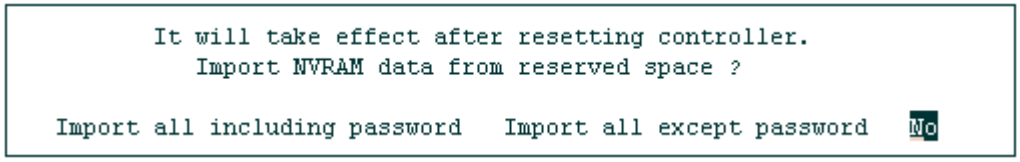
It will take effect after resetting controller.
Import NVRAM data from reserved space ?
Import all including password
Import all except password
No

4.7.7 Clearing Core Dump

Go to: System Functions > Controller Maintenance > Clear Core Dump

NOTE Upon seeing core dump events, power down and reboot your system after checking system events and correcting system faults. It is highly recommended to contact technical support immediately.

Please DO NOT clear the core dump data before causes of failures can be verified and corrected.

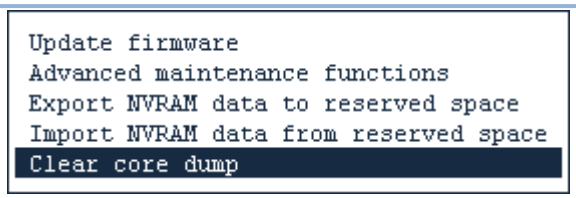
Update firmware
Advanced maintenance functions
Export NVRAM data to reserved space
Import NVRAM data from reserved space
Clear core dump

The Core Dump is a last resort option that helps debug critical issues in the event of serious system faults. When system firmware detects critical errors (such as multi-bit errors, PCI Bus Parity errors, etc.), it distributes configuration and error codes in cache memory into a core file in the 256MB disk reserved space. Galaxy's engineers can refer to these error codes from the core file conducted onto drive media if system finally crashes.

If system is recovered from serious faults later, you can execute the Clear Core Dump function to release disk space.

4.7.8 Adjusting the LCD Contrast

The controller LCD contrast is set at the factory to a level that should be generally acceptable. The controller is equipped with an LCD contrast adjustment circuit in case the factory-preset level needs to be adjusted either via the RS-232 terminal emulation menus or using the LCD keypad panel.

1. From the main menu, choose "View and Edit Peripheral Dev."

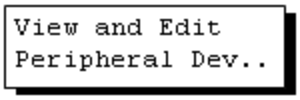

2. Press ENT on it, press arrow keys to scroll down, and select "Adjust LCD Contrast," press ENT to proceed, and then use the arrow keys to find an optimal setting.

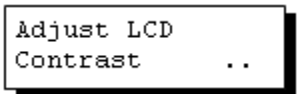

3. Press ESC to return to the previous menu.

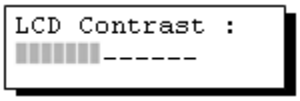

4.8 Configuring Controller Parameters

4.8.1 Changing the Controller Name

The controller name represents a RAID subsystem in a deployment that consists of numerous RAID subsystems. With dual-controller configurations, only one controller name is applied and will pass down to the surviving controller in the event of single controller failure.

LCD Keypad

1. Select “View and Edit Config Parms” from the Main Menu.

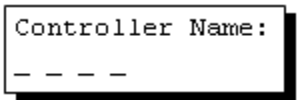

2. Choose “View and Edit Configuration Parameters,” “Controller Parameters," then press ENT.

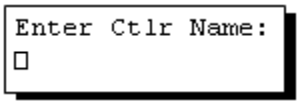

3. The current name will be displayed. Press ENT for two seconds and enter the new controller name by using the up or down arrow keys. Press ENT to move to another character and then press ENT for two seconds on the last digit of the controller name to complete the process.

Terminal

Go to: View and Edit Configuration Parameters > Controller Parameters > Controller Name

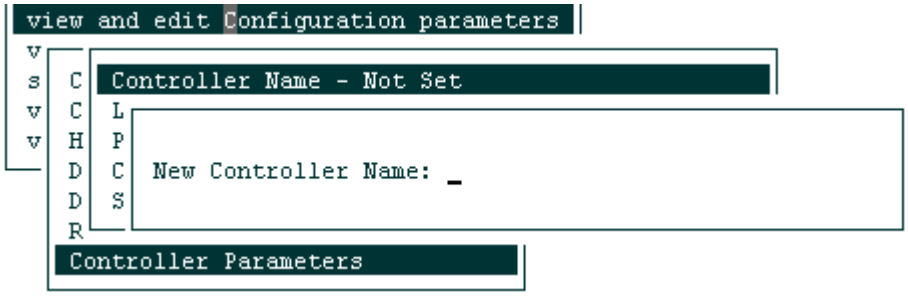

4.8.2 Showing the Controller Name

LCD Keypad

1. Choose “View and Edit Configuration Parameters,” “Controller Parameters,” then press ENT.

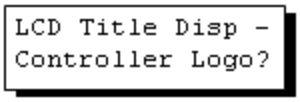

2. Use the up or down arrow keys to choose to display the embedded controller logo or any given name on the LCD initial screen.

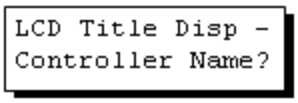

Terminal

Go to: View and Edit Configuration Parameters > Controller Parameters > LCD Title Display

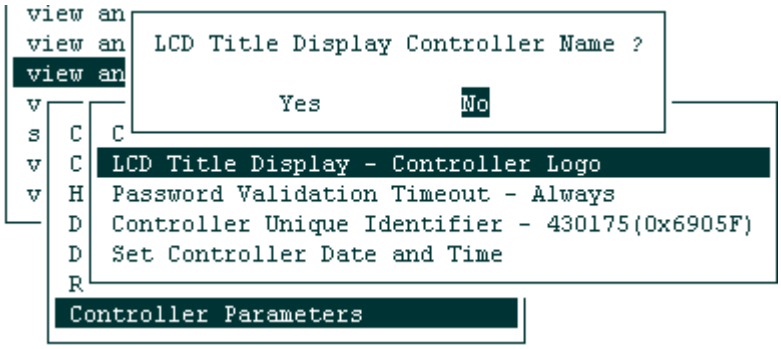

Choose to display the embedded controller model name or any given name on the LCD. Giving a specific name to each controller will make them easier to identify if you have multiple RAID systems that are monitored from a remote station.

4.8.3 Setting Password Validation Timeout

NOTE

The Always Check timeout will disable any attempts to make configuration changes without entering the correct password.

LCD Keypad

1. Choose “View and Edit Configuration Parameters,” “Controller Parameters,” then press ENT.

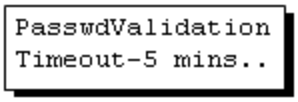

2. Select “Password Validation Timeout,” and press ENT. Press the up or down arrow keys to choose to enable a validation timeout from one to five minutes, or to “Always Check.”

Terminal

Go to: View and Edit Configuration Parameters > Controller Parameters > Password Validation Timeout

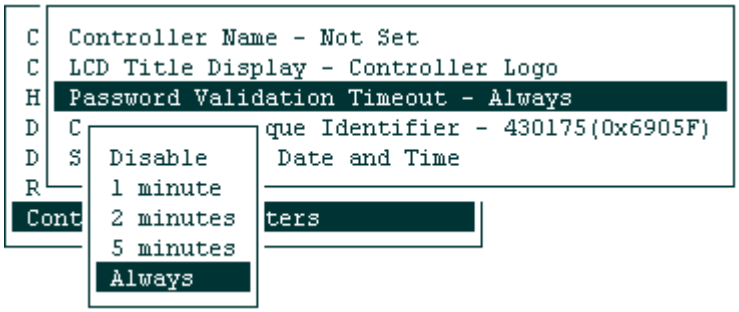

4.8.4 Setting a Unique Controller Identifier

What is the Controller Unique Identifier?

- A specific identifier helps RAID controllers to identify their counterpart in a dual-active configuration.
- The unique ID is generated into a Fibre Channel WWN node name for RAID controllers or RAID subsystems using Fibre Channel host ports. The node name prevents host computers from misaddressing the storage system during the controller failover/failback process in the event of single controller failure.
- The unique ID is also generated into a MAC address for the controller’s Ethernet port. The MAC address will be taken over by a surviving controller in the event of single RAID controller failure.

- When a controller fails and a replacement is combined as the secondary controller, the FC port node names and port names will be passed down to the replacement controller. The host will not acknowledge any differences so that controller failback is totally transparent.

LCD Keypad

3. Choose “View and Edit Configuration Parameters,” “Controller Parameters," then press ENT.

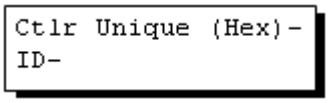

4. Press the up or down arrow keys to select “Ctlr Unique ID-,” then press ENT.

5. Enter any hex number between “0” and “FFFFF” and press ENT to proceed.

NOTE

Usually every RAID subsystem/controller comes with a default ID. In rare occasions should this identifier be changed.

There are chances that if you move a controller from a similar RAID system to another, that controller might have already acquired a unique ID from the original system’s EEPROM. As the result, its Fibre Channel port names can be identical to those on the system where the controller comes from. SAN port name conflicts can occur.

Terminal

Go to: View and Edit Configuration Parameters > Controller Parameters > Controller Unique Identifier

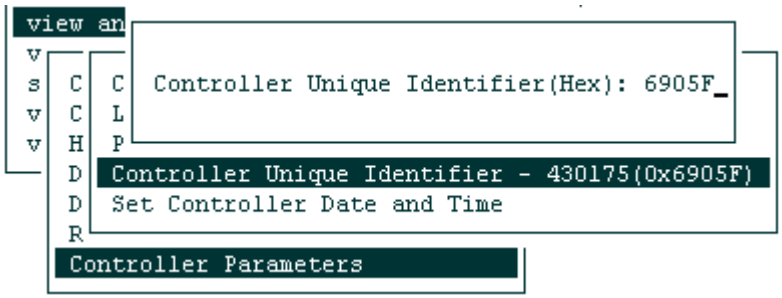

Enter any hex number between “0” and “FFFFF” for the unique identifier. The value you enter MUST be different for each controller.

Every Galaxy HDX4 subsystem comes with a default ID. This ID should be sufficient for avoiding WWNN and WWPN conflicts.

4.9 Installing RitePath Driver

Install Driver on the Application Server (In this case we are using Windows)

1. Select and execute the appropriate RitePath driver for your OS by a double-click. RitePath is included in your product CD, and the driver revisions can be acquired via technical support.

2. The progress indicator and a DOS prompt will appear.

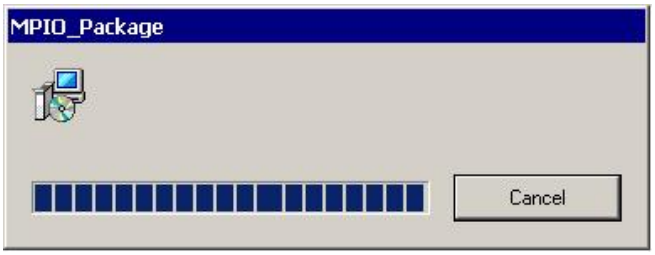

3. Press Y to confirm the legal notice.

4. Press Enter when the installation process is completed.
   Reboot your server for the configuration to take effect

5. You can check the availability of RitePath service using the Computer Management utility by a right-click on the My Computer icon. Select Manage, and select Services from the item tree.

6. Since the multi-pathing driver is already working, you can see multi-path device in Device Manager -> Disk Drives.

7. Upon seeing the Multi-Path Disk Device, you can start using the storage volumes from the redundant-controller iSCSI system.

5 Galaxy HDX4 Architecture

Before creating storage configurations, it is recommended you read through this chapter to gain understanding of the HDX4 series storage architecture. Comparisons among the standard earlier Galaxy models and the latest Galaxy HDX4 series are listed.

5.1 Differences from Other Galaxy Storage Series

5.1.1 List of Key Differences

The units presented to host as LUNs are different:

	HDX4 [DS] *	HDX, HDX2, HDX3
Mappable units (LUNs)	Partitions within Logical Volumes	Partitions within Logical Drives ***OR*** Partitions within Logical Volumes
Configuration interface	Storage Manager, LCD, terminal	Storage Manager, LCD, terminal

* Logical Drives are not mappable in the Galaxy HDX4 series.

Ways to present storage volumes are different:

HDX4 [DS]	HDX, HDX2, HDX3
Physical drives -> Logical Drives -> Logical Volumes -> Logical Partitions -> LUNs	Physical drives -> Logical Drives -> Logical Partitions -> LUNs Physical drives -> Logical Drives -> Logical Volumes -> Logical Partitions -> LUNs

5.1.2

Storage Components

Galaxy Series

Physical drives are included in logical drives. If a logical drive is not partitioned, all its capacity will appear as a single “partition 0.”

In a standard Galaxy, logical volumes and multiple partitions in an LD are optional.

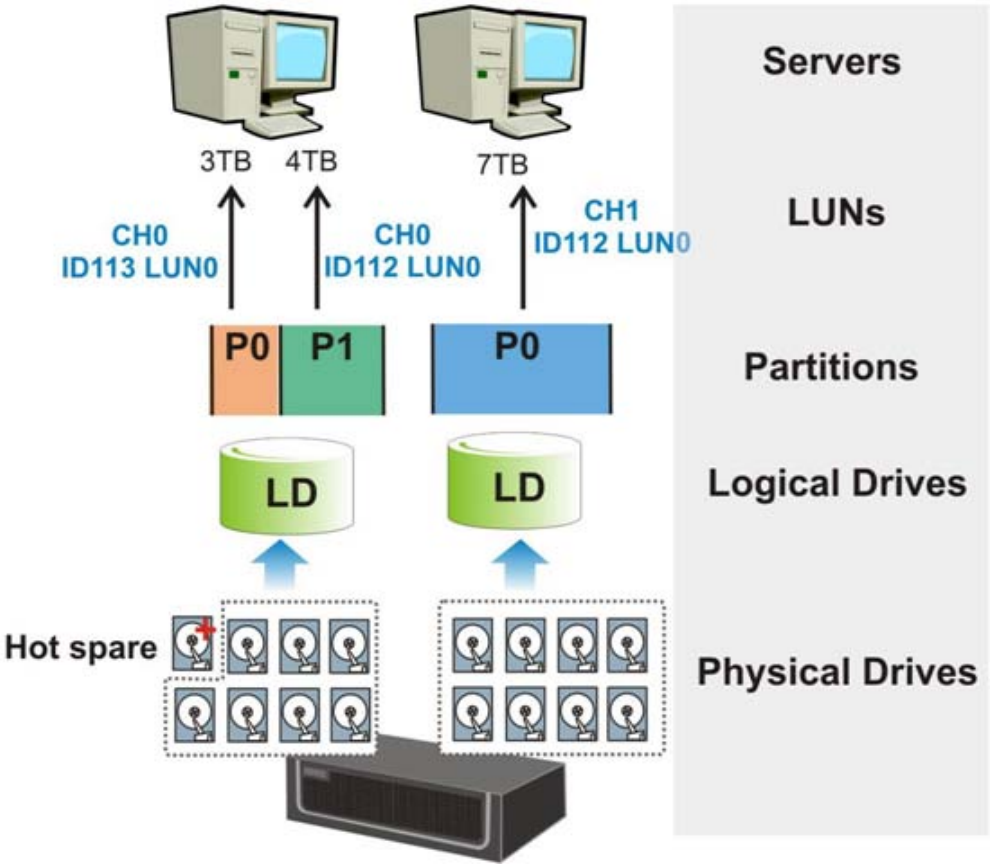

5.1.3 Data Services

The Data Service features include: Snapshot, Volume Copy, Volume Mirror, Thin Provision (from firmware version 386), and a scheduler tool for these features. Below is the availability of data services for different series:

	Galaxy HDX4	**HDX, HDX2, HDX3**
Data services	Yes	No
Scale-out & Load Balance	No	No
Virtualization & thin-provision	No	No

The Data Service functions in the HDX4 series can only be configured through the Galaxy Array Manager console. You can not create Data Service configurations using a RS232 terminal and LCD keypad.

5.1.4 Limitations

There are limitations for standard Galaxy HDX,HDX2, HDX3 and Galaxy HDX4 series:

- A standard Galaxy cannot be upgraded to Galaxy HDX4.
- You can not expand a logical volume by adding new members in the HDX4.
- A Galaxy HDX4 cannot be downgraded to a older standard Galaxy.
- A Physical LD configured in Galaxy cannot be moved and installed into a Logical Volume in Galaxy HDX4.
- You cannot perform Volume Copy or Volume Mirroring between an Galaxy HDX4 and a standard Galaxy.

Software features are separately-purchased and available via software licenses:

They include:

- Galaxy HDX4 Advanced In-System Replication (Snapshot + Volume Copy/Mirror of volumes within the same storage configuration consisting of 1 RAID system and multiple JBODs)
- Galaxy HDX4 Remote Replication (Copy or Mirror between volumes managed

by two different systems)

- Thin Provision features is supported from firmware version 386.

5.2 Samples of HDX4 Deployment

Before you start to configure a RAID system, make sure that hardware installation is completed before any configuration takes place.

5.2.1 Typical HDX4 Deployment

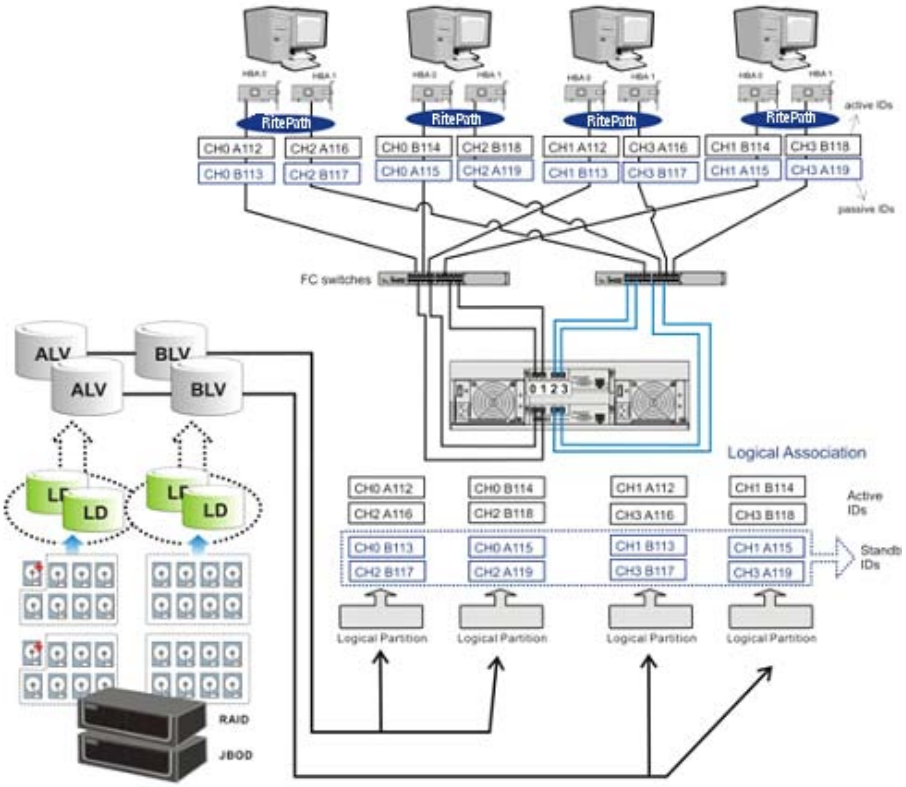

Shown above is a typical 8-port, redundant-controller system using AAPP (active-active-passive-passive) mapping. Doing so allows a LUN to be presented to host via multiple data paths to withstand cable disconnections or hardware failure. The passive paths do not carry data traffic in normal conditions, and will become active when active paths fail.

Port binding, zoning, file-locking, and other access control mechanisms should be implemented to avoid multiple servers from accessing the same storage volume.

Using a redundant-controller HDX4, system resource is manually separated by assigning logical volumes to different RAID controllers.

This page left blank intentionally

6 Creating Arrays and Mapping (LCD Panel)

A navigation roadmap for the configuration menu options through LCD keypad is separately available as a PDF file. You may check your Product Utility CD or contact technical support for the latest update.

Before you start to configure a RAID system, make sure that hardware installation is completed before any configuration takes place. Power on your RAID system.

Notes of Power up:

- If your Galaxy HDX4 RAID system comes with dual-redundant RAID controllers, your system's LCD panel can provide access to the operating status screen of the Secondary controller. However, in a dual-controller configuration, only the Primary controller responds to user's configuration.
- Each controller's operating mode is indicated by the flashing digit on the upper right of the LCD screen as "A" or "B." If the LCD displays "B," that means the LCD screen is currently displaying Slot B controller messages. Press both the Up and Down arrow keys for one second to switch around the access to different RAID controllers.

6.3 Working with a Logical Drive

6.3.1 Creating a Logical Drive

1. To create a logical drive, press ENT for two seconds to enter the Main Menu. Use the up or down arrow keys to navigate through the menus. Choose "View and Edit Logical Drives," and then press ENT.

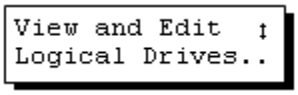

2. Press the up or down arrow keys to select a logical drive index entry, then press ENT for two seconds to proceed. "LD" is short for Logical Drive.

```
LD=0
Not Defined    ?
```

3. Use the up or down arrow keys to select the desired RAID level, then press ENT for two seconds. "TDRV" (Total Drives) refers to the number of all available disk drives.

```
TDRV=24 Create
LG Level=RAID6 ?
```

6.3.2 Choosing Member Drives

1. Press ENT for two seconds; the message, “RAID X selected To Select drives”, will prompt. Confirm your selection by pressing ENT.

```
RAID 6  Selected
To Select drives
```

2. Press ENT, then use the up or down arrow keys to browse through the available drives.

```
Slot=* 239367MB
New DRV SEAGATE
```

3. Press ENT again to select/deselect individual disk drives. An asterisk (*) mark will appear on the selected drive(s). To deselect a drive, press ENT again on the selected drive. The (*) mark will disappear.

```
SLOT=1 239367MB
*LG=0 SL
```

4. After all the desired hard drives have been selected, press ENT for two seconds to continue.

6.3.3 Setting Maximum Drive Capacity

1. You may enter the following screen to “Change Logical Drive Parameter” by pressing ENT before initializing the logical drive.

```
Change Logical
Drive Parameter?
```

2. Choose “Maximum Drive Capacity,” then press ENT. The maximum drive capacity refers to the maximum capacity that will be used in each individual member drive.

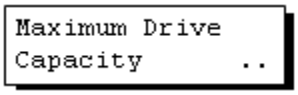

3. If necessary, use the up and down arrow keys to change the maximum size that will be used on each drive.

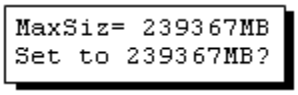

6.3.4 Assigning a Spare Drive

1. You may enter the following screen to “Change Logical Drive Parameter” by pressing ENT before initializing the logical drive.

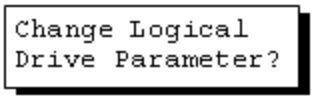

2. Press the up or down arrow keys to choose “Spare Drive Assignments,” then press ENT.

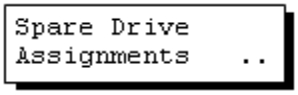

3. Available disk drives will be listed. Use the up or down arrow keys to browse through the drive list, then press ENT to select the drive you wish to use as the Local (Dedicated) Spare Drive.

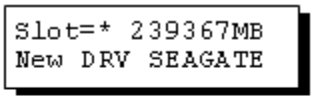

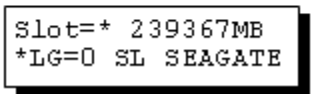

4. Press ENT again for two seconds.

6.3.5 Viewing Reserved Disk Space

This menu allows you to see the size of disk reserved space. Default is 256MB. The reserved space is used for storing array configuration and other non-volatile information.

6.3.6 Setting Write Policy

This menu allows you to set the caching mode policy for this specific logical drive. “Default” is a neutral value that is coordinated with the subsystem’s general caching mode setting. Other choices are “Write-back” and “Write-through.”

1. You may enter the following screen to “Change Logical Drive Parameter” by pressing ENT before initializing the logical drive.

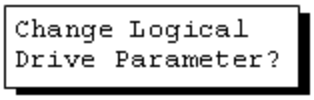

2. Press ENT once to change the status digits into a question mark “?”.

3. Use the arrow keys to select “Default,” “Write-back,” or “Write-through.”

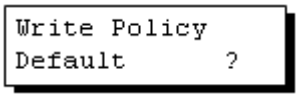

4. Press ENT for two seconds to confirm your change.

NOTE

The “Write-back” and “Write-through” parameters are permanent for specific logical drives. The “Default” selection, however, is more complicated and more likely equal to “not specified.”

If set to “Default,” a logical drive’s write policy is determined not only by the system’s general caching mode setting, but also by the “Event trigger” mechanisms. The “Event Trigger” mechanisms automatically disable the write-back caching and adopt the conservative “Write-through” mode in the event of battery or component failures.

6.3.7 Setting Initialization Mode

This menu allows you to determine if the logical drive is immediately accessible. If the Online method is used, data can be written onto it before the array’s initialization is completed. You may continue with other array configuration processes, e.g., including this array in a logical volume.

Array initialization can take a long time especially for those comprising a large capacity and parity data. Setting to “Online” means the array is immediately accessible and that the controller will complete the initialization in the background or I/Os become less intensive.

1. You may enter the following screen to “Change Logical Drive Parameter” by pressing ENT before initializing the logical drive.

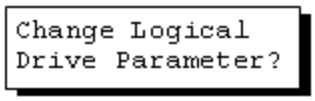

2. Press ENT once to change the status digits into a question mark “?”.

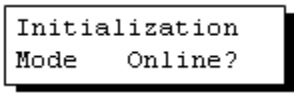

3. Use the arrow keys to select either the “Online” or the “Off-line” mode.

4. Press ENT for two seconds to confirm your change.

6.3.8 Setting Stripe Size

This menu allows you to change the array stripe size. Setting to an incongruous value can severely drag performance. This item should only be changed when you can test the combinations of different I/O sizes and array stripe sizes and can be sure of the performance gains it might bring you.

For example, if the I/O size is 256k, data blocks will be written to two of the member drives of a 4-drive array while the RAID firmware will read the remaining member(s) in order to generate the parity data.

* For simplicity reasons, we use RAID3 in the samples below.

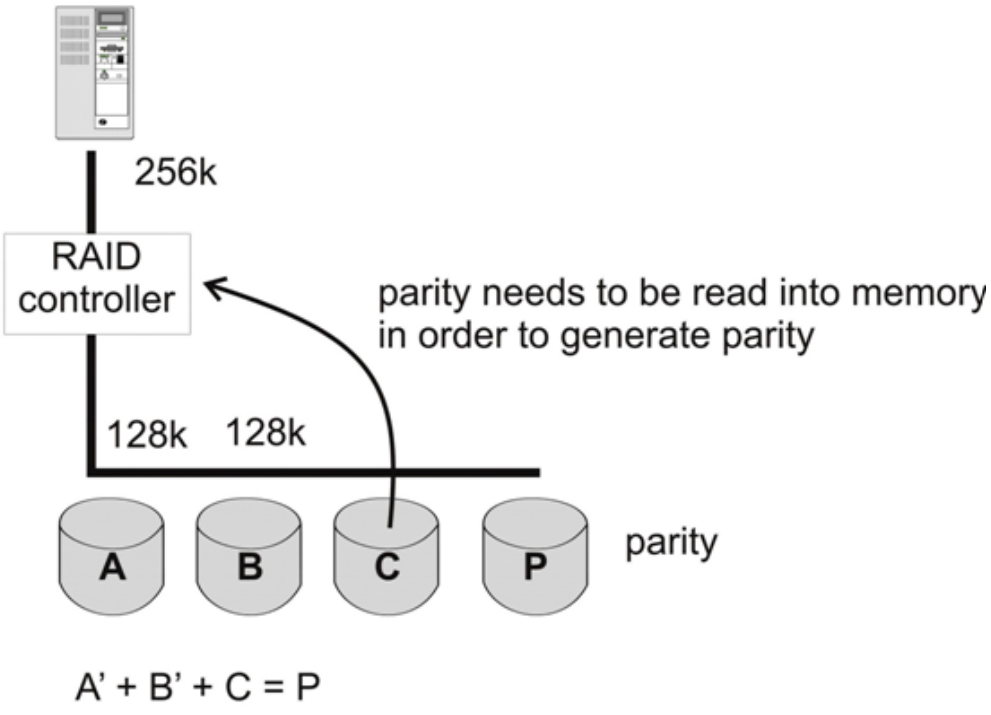

In an ideal situation, a 384k I/O size allows data to be written to 3 member drives and parity data to be simultaneously generated without the effort to consult data from other members of an array.

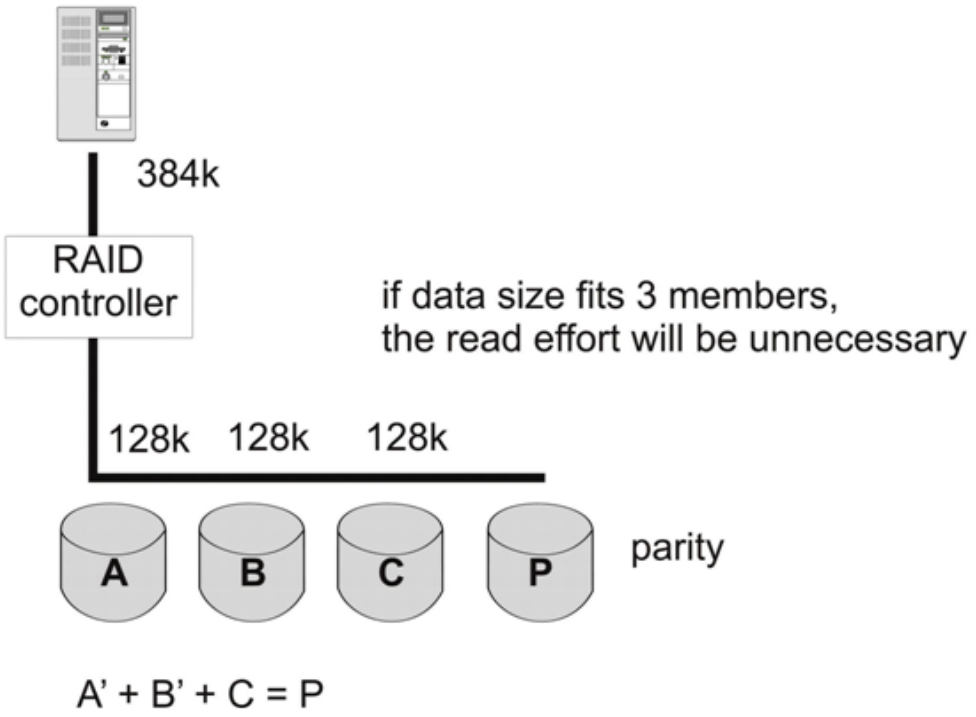

If the I/O size is larger than the combined stripe depths, the extra data blocks will be written to the member drives on the successive spins, and the read efforts will also be necessary for generating parity data.

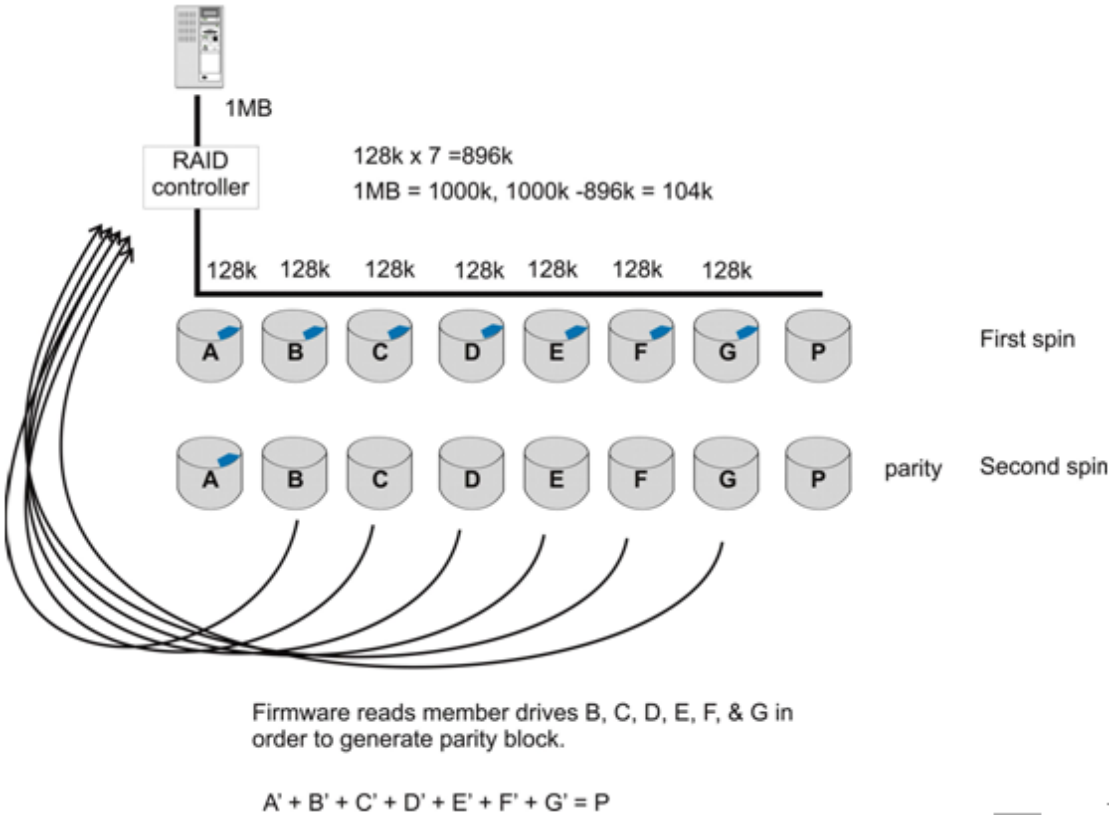
1MB
RAID controller
128k x 7 =896k
1MB = 1000k, 1000k -896k = 104k
128k 128k 128k 128k 128k 128k 128k
A B C D E F G P
First spin
A B C D E F G P
parity
Second spin
Firmware reads member drives B, C, D, E, F, & G in order to generate parity block.
A' + B' + C' + D' + E' + F' + G' = P

Although the real-world I/Os do not always perfectly fit the array stripe size, matching the array stripe size to your I/O characteristics can eliminate drags on performance (hard drive seek and rotation efforts) and will ensure the optimal performance.

Listed below are the default values for different RAID levels.

RAID level	Stripe Size
RAID0	128KB
RAID1	128KB
RAID3	16KB
RAID5	128KB
RAID6	128KB
NRAID	128KB

1. You may enter the following screen to “Change Logical Drive Parameter” by pressing ENT before initializing the logical drive.

```
Change Logical
Drive Parameter?
```

2. Press ENT once to change the status digits into a question mark “?”.

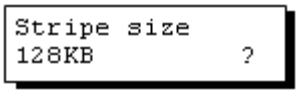

3. Use the arrow keys to select a desired stripe size.

4. Press ENT for two seconds to confirm your change.

6.3.9 Initializing a Logical Drive

1. Press ESC to return to the previous menu. Use the up or down arrow keys to select “Create Logical Drive?”

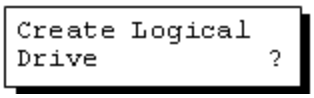

2. Press ENT for two seconds to start initializing the logical drive.

3. Repeat the above processes to create more logical drives. These logical drives will form logical volumes later.

The Online Mode:

If the online initialization method is applied, the array will be immediately available for use. The array initialization runs in the background and the array is immediately ready for I/Os and further configurations. Engineers can continue configuring the RAID subsystem.

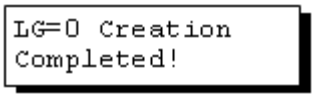

The Offline Mode:

The RAID controller will immediately start to initialize the array parity if the “offline” mode is applied. Note that if NRAID or RAID0 is selected, initialization time is short and completes almost within a second.

The logical drive’s information displays when the initialization process is completed. If the “online” mode is adopted, array information will be displayed immediately.

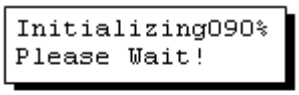

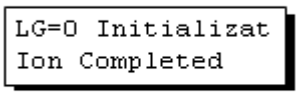

NOTE

Due to the operation complexity, the RAID Migration option is not available using the LCD keypad panel.

6.3.10 Naming a Logical Drive

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Drives..," then press ENT.

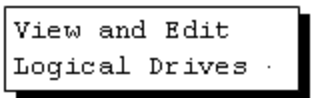

2. Press the up or down arrow keys to select a logical drive, then press ENT.

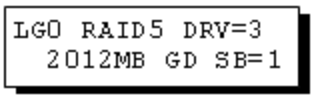

3. Press the up or down arrow keys to select “Logical Drive Name," then press ENT.

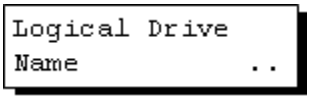

4. Press the up or down arrow keys to change the character of the flashing cursor. Press ENT to move the cursor to the next space. The maximum number of characters for a logical drive name is 32.

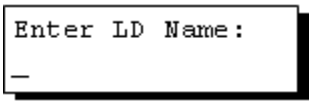

A similar option is also found with the logical volumes.

6.3.11 Deleting a Logical Drive

NOTE

Deleting a logical drive erases all data stored in it.

You can not delete a logical drive if it has been included in a logical volume.

You can not disband members of a logical volume unless all its host LUN mappings are erased.

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Drives," then press ENT.

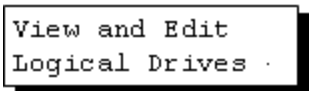

2. Press the up or down arrow keys to select a logical drive, then press ENT.

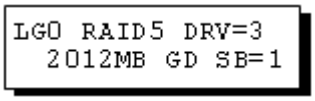

3. Use the up or down arrow keys to select “Delete Logical Drive," then press ENT.

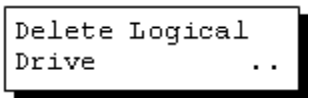

4. Press ENT for two seconds to confirm.

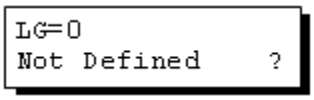

6.3.12 Deleting the Partition of a Logical Drive

NOTE

Whenever there is a partition change, data will be erased. Prior to partition change, you have to remove its associated host LUN mappings. After the partition change, you also need to re-arrange the disk volumes from your host system OS.

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Volumes..,” then press ENT.

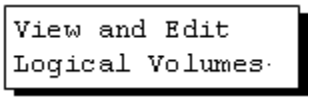

2. Press the up or down arrow keys to select a logical drive, then press ENT.

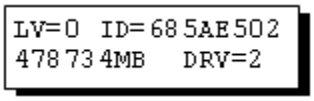

3. Press the up or down arrow keys to choose “Partition Logical Volume," then

press ENT.

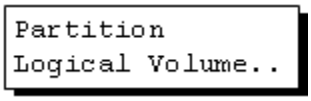

4. The first partition's information will be shown on the LCD. Press the up or down arrow keys to browse through the existing partitions in the logical volume. Select a partition by pressing ENT for two seconds.

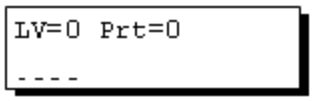

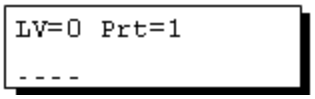

5. Use the up or down arrow keys to change the number of the flashing digit to "0," then press ENT to move to the next digit. After changing all the digits, press ENT for two seconds.

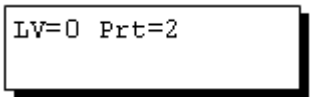

The disk space of the deleted partition will be automatically allocated to the free space partition as diagrammed below. For example, if partition 1 is deleted, its disk space will be added to free space.

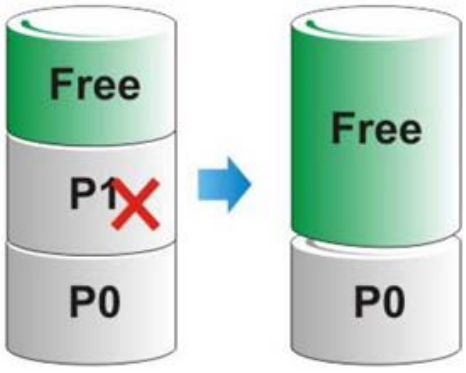

6.4 Working with Logical Volumes

NOTE

Logical volume is a must for RAID configuration in a HDX4 series.

6.4.1 Creating a Logical Volume

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Volume," then press ENT.

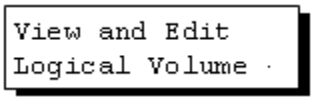

2. Press the up or down arrow keys to select an undefined index entry for a logical volume, then press ENT for two seconds to proceed. "LV" is short for Logical Volume.

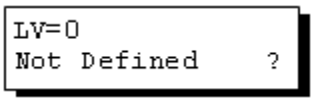

3. Proceed to select one or more logical drives as the members of a logical volume. Press ENT to proceed. “LD” is short for Logical Drive.

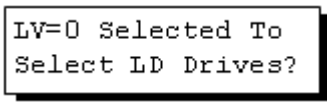

4. Use the up or down arrow keys to browse through the logical drives.

5. Press ENT again to select/deselect the members. An asterisk (*) mark will appear in front of a selected logical drive.

6. After all the desired logical drive(s) have been selected, press ENT for two seconds to continue.

7. Two sub-menus will appear.

6.4.2 Setting the Initialization Mode

Array initialization can take a long time especially for those comprised of a large capacity. Setting to “Online” means the array is immediately accessible and that the controller will complete the initialization in the background and when I/O demands become less intensive. Default is Online.

1. Press ENT once to change the status digits into a question mark “?”.

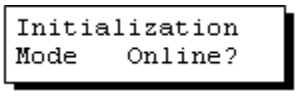

2. Use the arrow keys to select either the “Online” or the “Off-line” mode.

3. Press ENT for two seconds to confirm your change.

6.4.3 Setting the Write Policy

NOTE

The “Write-back” and “Write-through” parameters are permanent for specific logical drives. The “Default” selection, however, is more complicated and more likely equal to “not specified.”

If set to “Default,” a logical drive’s write policy is controlled not only by the subsystem’s general caching mode setting, but also by the “Event trigger” mechanisms. The “Event Trigger” mechanisms automatically disable the write-back caching and adopt the conservative “Write-through” mode in the event of a battery or component abnormalities.

This menu allows you to set the caching mode policy for this specific logical volume. “Default” is a neutral value that is coordinated with the controller’s general caching mode setting. Other choices are “Write-back” and “Write-through.”

1. Press ENT once to change the status digits into a question mark “?”

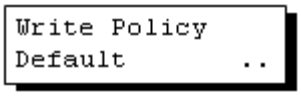

2. Use the arrow keys to select “Default,” “Write-back,” or “Write-through.”

3. Press ENT for two seconds to confirm your change.

4. When you are finished setting the preferences, press ENT for two seconds to display the confirm box. Press ENT for two seconds to start initializing the logical volume.

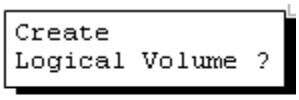

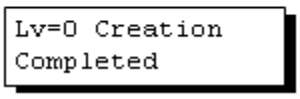

5. A message shows that the logical volume has been successfully created.

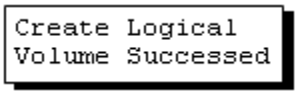

6. Press ESC to clear the message.

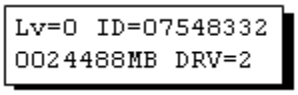

7. Logical volume information will be displayed next.

> NOTE
>
> Once a logical drive is included in a logical volume, its “Controller Assignment” option will disappear. The controller assignment option displays under the logical volume sub-menu instead.

6.4.4 Assigning a Logical Volume (Dual-active Controllers)

In a dual-controller configuration, you may choose to assign this logical volume to the Slot B controller (Default is Slot A, the default dominant/master controller). The assignment can take place during or after the initial configuration.

In a dual-controller configuration, the assignment menus should appear as listed on the right.

1. Press ENT on a configured logical volume. Use arrow keys to select “Logical Volume Assignment..”, and press ENT to proceed. Press ENT for two seconds to confirm.

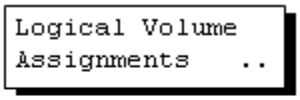

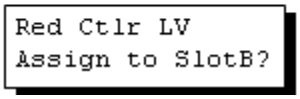

2. Press ESC, and the LCD will display the logical volume’s information when initialization is completed.

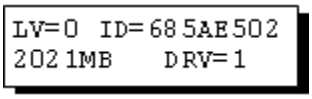

In theory, you should divide the workload on a system by assigning half of logical volumes to Slot A controller, and another half to the Slot B controller.

6.4.5 Partitioning a Logical Volume

NOTE

Partitioning is a requirement for building storage volumes in a HDX4 series system.

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Volume," then press ENT.

```
View and Edit
Logical Volume
```

2. Use the up or down arrow keys to select a logical volume, then press ENT.

```
LV=0 ID=685AE502
2021MB   DRV=1
```

3. Use the up or down arrow keys to select “Partition Logical Volume,” then press ENT.

```
Partition
Logical Volume..
```

4. The total capacity of the logical volume will be displayed as one partition. Press ENT for two seconds to change the size of the first partition.

```
LV=0   Part=0:
   2021MB
```

5. Use the up or down arrow keys to change the number of the flashing digit, (see the arrow mark) then press ENT to move to the next digit.

```
LV=0 Part=0:
  700MB       ?
```

6. After changing all the digits, press ENT for two seconds to confirm the capacity of this partition. You may then use arrow keys to move to the next partition to configure more logical partitions.

```
LV=0 Partition=1
  1321MB       ?
```

The rest of the drive space will be automatically allocated to the next partition. You may repeat the process to create up to 64 partitions using the same method described above.

7. Press ESC several times to return to the Main Menu.

NOTE

If operating with a Unix-based system, reset the system for the configuration to take effect if any changes have been made to partition sizes and partition arrangement.

6.4.6 Mapping Logical Partitions Drive to Host LUN

NOTE

The current firmware revisions support the cross-controller ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs. However, mapping to both controllers' IDs is usually beneficial when it is difficult making the fault-tolerant host links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. The cross-controller mapping also makes sense in a clustered server environment. Currently, external SAS switches are not popular on the market. For Fibre-host systems, fault-tolerant links can easily be made with the help of external bypass such as Fibre Channel switches. For details of fault-tolerant link connections, please refer to your system Hardware Manual.

The idea of host LUN mapping is diagrammed as follows:

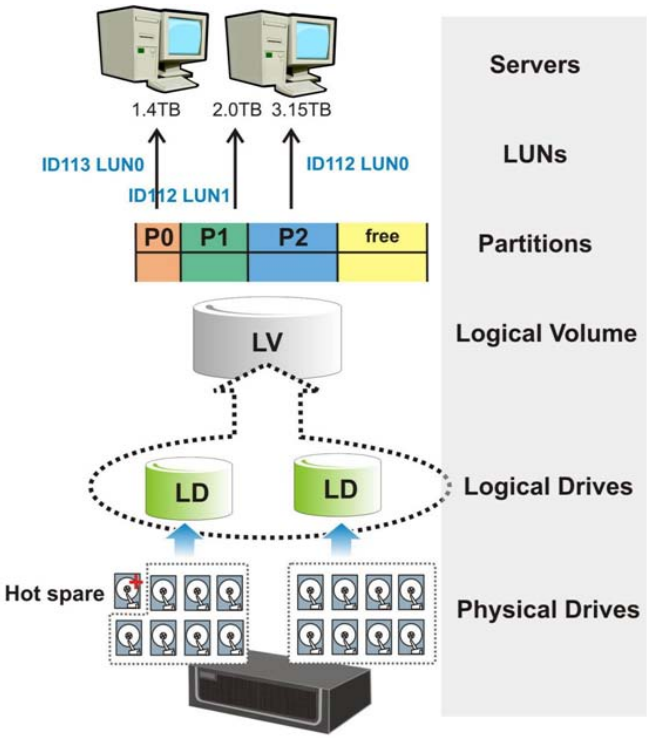

NOTE

Your subsystem comes with one Slot A ID and one Slot B ID only. If you need more host channel IDs, you need to manually create them. Please enter "View and Edit Channels" menu to create or remove a host ID.

1. The first available ID on the first host channel appears (usually channel0).

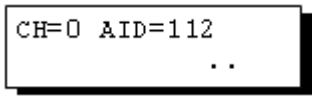

2. Press the up or down arrow keys to select an existing host ID, and then press ENT for two seconds to confirm.

3. Press the up or down arrow keys to select the type of logical configuration to be associated with a host ID/LUN. "Map to Logical Volume".

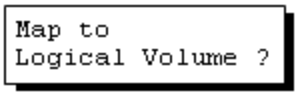

4. Confirm your choice by pressing ENT for two seconds.

5. Press the up or down arrow keys to select a LUN number under host ID, then press ENT to proceed.

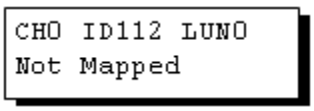

6. Press ENT for two seconds to confirm the selected LUN mapping.

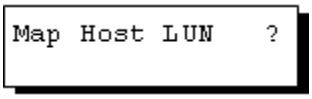

7. Press the up or down arrow keys to select a logical volume or a partition within.

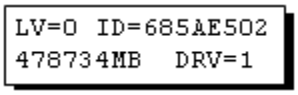

8. Press ENT for two seconds to map the selected partition to this LUN.

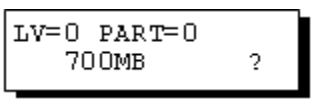

9. Press ENT for two seconds when prompted by "Map Host LUN" to proceed.

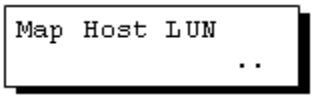

10. Mapping information will be displayed on the subsequent screen. Press ENT for

two seconds to confirm the LUN mapping.

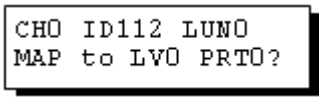

11. The mapping information will appear for the second time. Press ENT or ESC to confirm, and the host ID/LUN screen will appear.

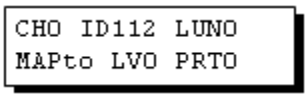

12. Use the arrow keys to select another ID or LUN number to continue mapping other logical configurations or press ESC for several times to leave the configuration menu.

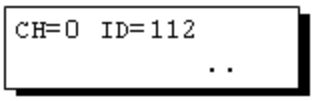

When any of the host ID/LUNs is successfully associated with a logical array, the “No Host LUN” message in the initial screen will change to “Ready.”

6.4.7 Deleting Host LUNs

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Host Luns", then press ENT.

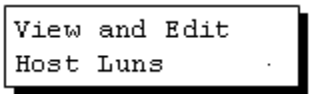

2. Press the up or down arrow keys to select a host ID, then press ENT to proceed.

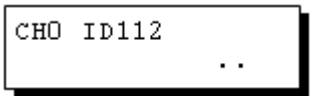

3. Use the up or down arrow keys to browse through the LUN number and its LUN mapping information.

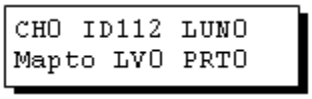

4. Press ENT on the LUN you wish to delete.

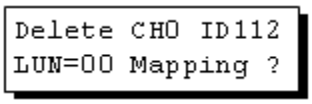

5. Press ENT for two seconds to confirm deletion. The deleted LUN has now been unmapped.

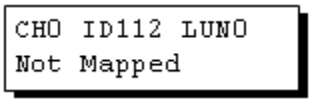

6.5 Assigning Spare Drive and Rebuild Settings

6.5.1 Adding a Local Spare Drive

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press ENT.

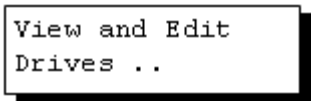

2. Disk drive information will be displayed on the LCD. Press the up or down arrow keys to select a drive that is stated as “NEW DRV” or “USED DRV” that has not been included in any logical drive, nor specified as a “FAILED” drive, then press ENT to select it.

3. Press the up or down arrow keys to select “Add Local Spare Drive,” then press ENT.

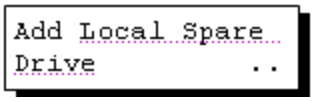

4. Press the up or down arrow keys to select the logical drive where the Local Spare Drive will be assigned, then press ENT for two seconds to confirm.

5. The message “Add Local Spare Drive Successful” will be displayed on the LCD.

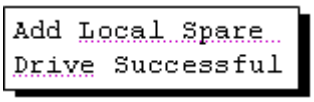

6.5.2 Adding a Global Spare Drive

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press ENT.

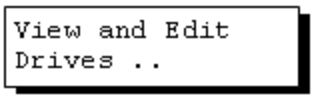

2. Disk drive information will be displayed on the LCD. Press the up or down arrow keys to select a disk drive that has not been assigned to any logical drive, then press ENT.

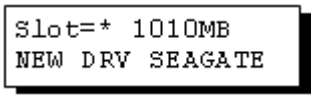

3. Press the up or down arrow keys to select “Add Global Spare Drive,” then press ENT

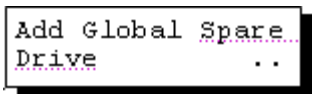

4. Press ENT again for two seconds to add the spare drive. The message, “Add Global Spare Drive Successful,” will be displayed on the screen.

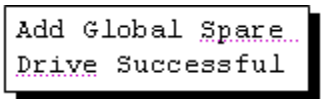

NOTE

Assigning a hot-spare to an array composed of drives of a different interface type should be avoided. For example, a SATA Global spare may accidentally participate in the rebuild of an array using SAS members. It is better to prevent mixing SAS and SATA drives in a logical drive configuration.

6.5.3 Adding an Enclosure Spare Drive

In environments where RAID volumes might span across several enclosures, e.g., using JBODs, this option can designate a spare drive to rebuild only a failed drive within the same enclosure. Using enclosure spares prevents a spare drive to join the rebuild in another enclosure.

1. To create an Enclosure Spare Drive, press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press ENT.

2. Disk drive information will be displayed on the LCD. Press the up or down arrow keys to select a disk drive that has not been assigned to any logical drive, then press ENT.

3. Press the up or down arrow keys to select “Add Enclosure Spare Drive,” then press ENT.

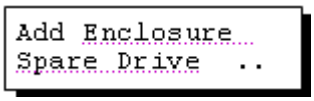

4. When the last digit changes to a question mark “?”, press ENT again for two seconds to create the enclosure spare. The message, “Add Spare Drive Successful,” will be displayed on the screen.

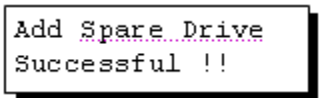

5. Press ESC and the drive status displays as shown on the right.

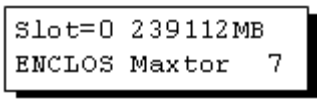

6.5.4 Deleting Spare Drive (Global / Local/Enclosure Spare Drive)

1. Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press ENT.

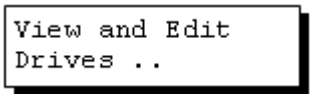

2. Drive information will be displayed on the LCD. Press the up or down arrow keys to select the spare drive you wish to delete, then press ENT.

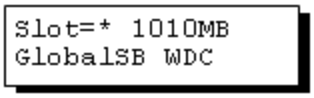

3. Press the up or down arrow keys to select “Delete Spare Drive," then press ENT to continue.

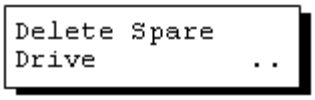

4. Press ENT for two seconds to delete the spare drive.

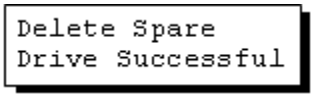

6.6 Restoring Firmware Default

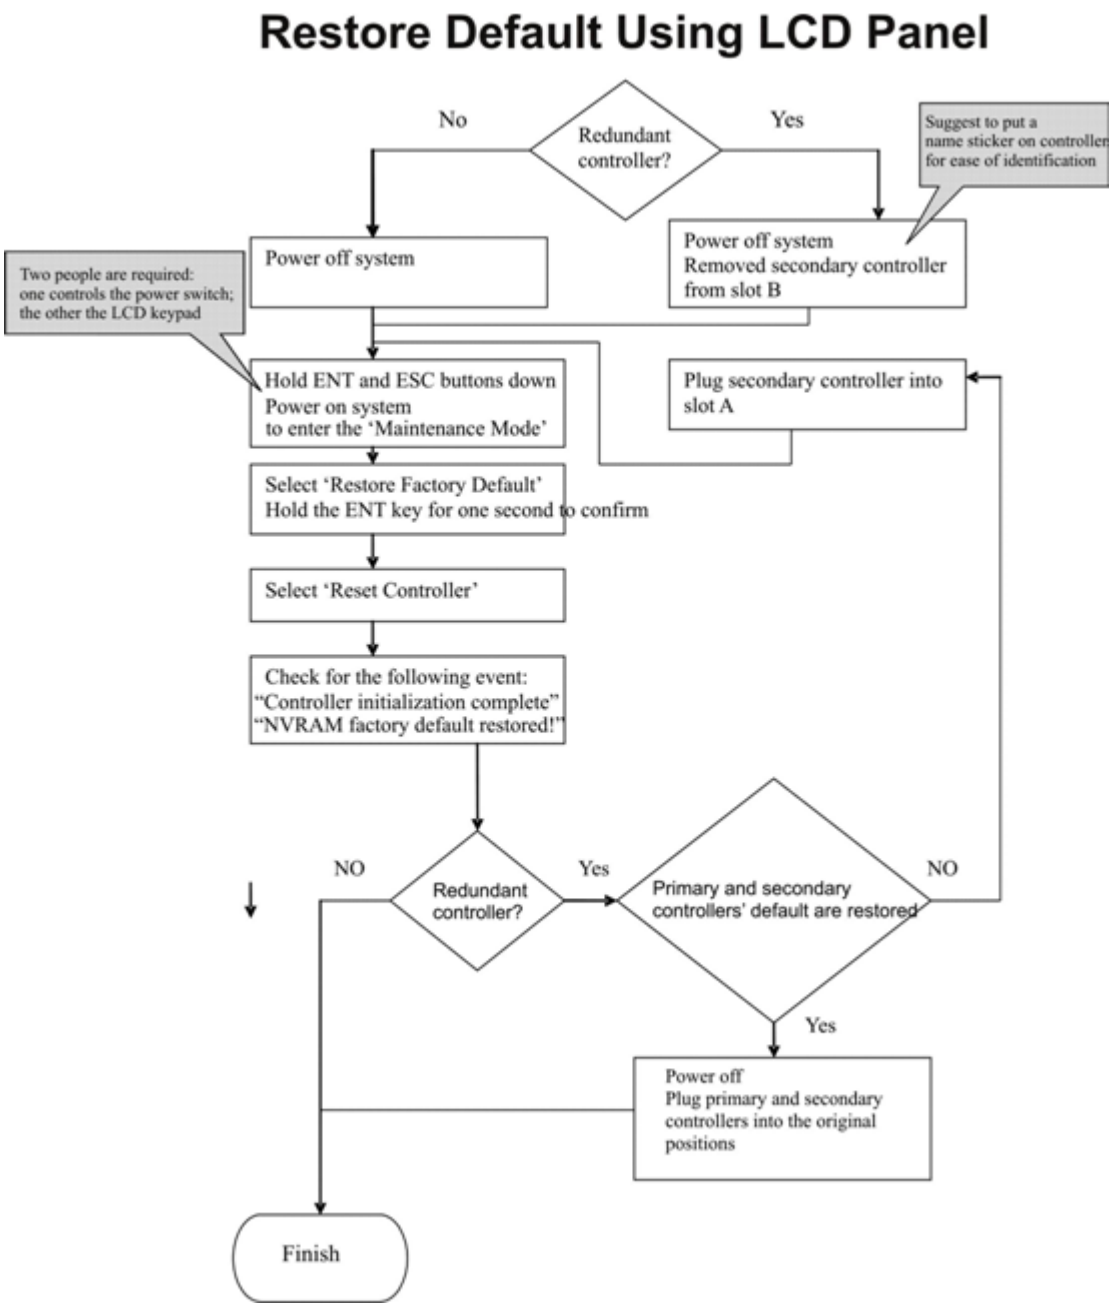
Restore Default Using LCD Panel
Redundant controller?
No
Yes
Suggest to put a name sticker on controller for ease of identification
Power off system
Power off system
Removed secondary controller from slot B
Two people are required: one controls the power switch; the other the LCD keypad
Hold ENT and ESC buttons down
Power on system
to enter the 'Maintenance Mode'
Plug secondary controller into slot A
Select 'Restore Factory Default'
Hold the ENT key for one second to confirm
Select 'Reset Controller'
Check for the following event:
"Controller initialization complete"
"NVRAM factory default restored!"
NO
Redundant controller?
Yes
Primary and secondary controllers' default are restored
NO
Yes
Power off
Plug primary and secondary controllers into the original positions
Finish

7 Creating Arrays and Mapping (Terminal)

Hardware installation should be completed before powering on your RAID enclosure. The subsystem and disk drives must be properly configured and initialized before the host computer can access the storage capacity. The text-based, menu-driven configuration and administration utility resides in the controller's firmware.

Open the initial terminal screen: use the arrow keys to move the cursor bar through the menu items, then press [ENTER] to select a terminal emulation mode, and [ESC] to dismiss current selection and/or to return to the previous menu/screen.

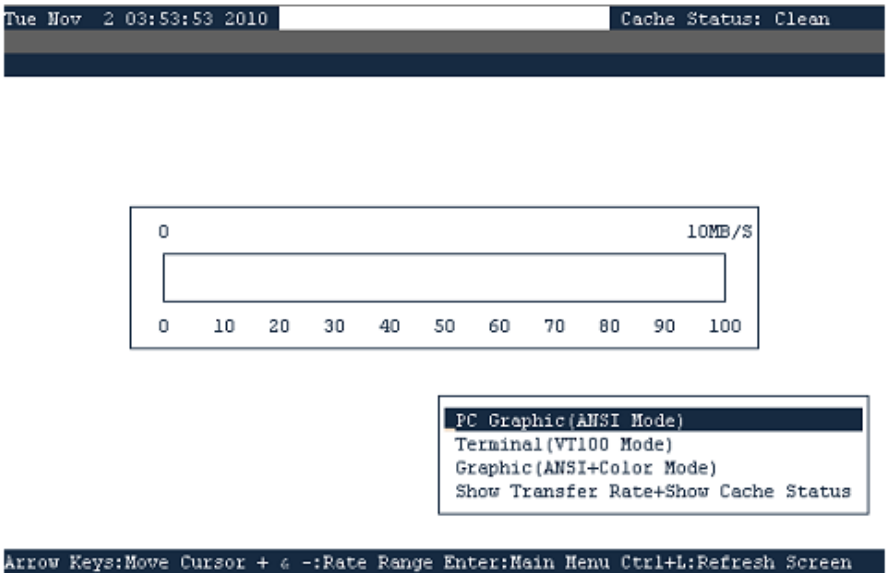
Tue Nov 2 03:53:53 2010
Cache Status: Clean
0
10MB/S
0 10 20 30 40 50 60 70 80 90 100
PC Graphic(ANSI Mode)
Terminal(VT100 Mode)
Graphic(ANSI+Color Mode)
Show Transfer Rate+Show Cache Status
Arrow Keys:Move Cursor + & -:Rate Range Enter:Main Menu Ctrl+L:Refresh Screen

7.1 Working with Physical Disks

Go to: View and Edit Drives

Prior to configuring individual disk drives into a logical drive, it is necessary to understand the status of all physical drives in your enclosure.

JBOD	Slot	ChNo	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	1	2	0	238214	300MB	0	ON-LINE	HDT722525DLA380
	2	2	1	238214	300MB	0	ON-LINE	HDT722525DLA380
	3	2	2	238214	300MB	GLOBAL	STAND-BY	HDT722525DLA380
	4	2	3		NONE		ABSENT	
	5	2	4		NONE		ABSENT	
	6	2	5		NONE		ABSENT	
	7	2	6		NONE		ABSENT	
	8	2	7		NONE		ABSENT	

Use the arrow keys to scroll down to “View and Edit Drives” to display information on all the physical drives installed.

Physical hard drives are listed in the “View and Edit Drives” table. Use the arrow keys to scroll the table. First examine whether there is any drive installed but not listed here. If a disk drive is installed but not listed, the drive may be faulty or not installed correctly. Reinstall the hard drives and contact your supplier for replacement drives.

NOTE

Drives of the same brand/model/capacity may not have the same block number.

The basic read/write unit of a hard drive is block. If members of a logical drive have different block numbers (capacity), the smallest block number will be taken as the maximum capacity to be used in every drive when composing a logical drive. Therefore, use drives of the same capacity.

You may assign a Spare Drive to a logical drive whose members have a block number equal or smaller than the Local/Global Spare Drive, but you should not do the reverse.

7.2 Working with Logical Drives

7.2.1 Creating a Logical Drive

Go to: View and Edit Logical Drives

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	0	C	#LN	#SB	#FL	NAME
A0	2B699693	0	RAID1	238214	GOOD			7	B	2	1	0	
1													
2													
3													
4													
5													
6													
7													

1. For the first logical drive on the RAID subsystem, simply choose the first logical drive index entry, “LG 0,” and press [ENTER] to proceed. You may create as many as 64 logical drives or more using drives in a RAID subsystem or in the expansion enclosures. The number of disk drives to be included in a logical drive depends on the capacity and performance concerns. The following are very rough examples using an 8-member RAID5:

- RAID5 LD capacity = [no. of HDDs -1(parity drive)] x single-drive capacity
  Exp. (8-1) x 1TB = 7TB
- LD performance: MB/s in pure reads [no. of HDDs - 1 (parity drive) x 100MB/s (15k SAS approx.)] x 85% (15% parity and I/Os handling overhead)
  Exp. (8-1) x 100 x 85% = 595 MB/s
- LD performance: random IOPS [no. of HDDs -1 (parity) x 180 IOPS (15k SAS approx.)] x 85% (15% parity and I/Os handling overhead)
  Exp. (8-1) x 180 x 85% = 1071 IOPS

2. File system caching, read/write characteristics, IO size, access patterns, and application buffering will all be concerns with the overall system performance.

3. When prompted to “Create Logical Drive?,” select Yes and press [ENTER] to proceed.

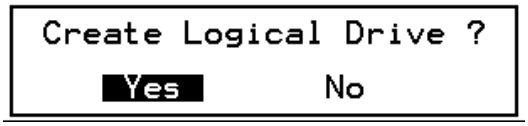

4. A pull-down list of supported RAID levels will appear. Choose a RAID level for this logical drive. In this chapter, RAID 6 will be used to demonstrate the configuration process.

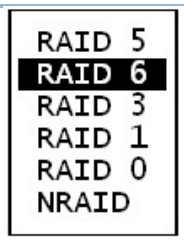

5. Choose your member drive(s) from the list of physical drives available. Tag the drives for inclusion by positioning the cursor bar on the drive and then pressing [ENTER]. An asterisk "*" mark will appear in front of the selected drive(s). To deselect the drive, press [ENTER] again on the selected drive and the asterisk "*" will disappear. Use the same method to select more member drives.

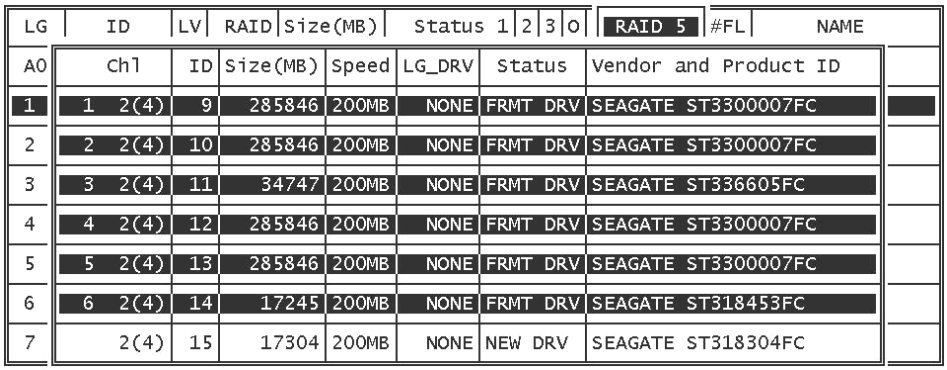

LG | ID | LV | RAID | Size(MB) | Status 1 2 3 0 | RAID 5 | #FL | NAME

A0	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	1 2(4)	9	285846	200MB	NONE	FRMT DRV	SEAGATE ST3300007FC
2	2 2(4)	10	285846	200MB	NONE	FRMT DRV	SEAGATE ST3300007FC
3	3 2(4)	11	34747	200MB	NONE	FRMT DRV	SEAGATE ST336605FC
4	4 2(4)	12	285846	200MB	NONE	FRMT DRV	SEAGATE ST3300007FC
5	5 2(4)	13	285846	200MB	NONE	FRMT DRV	SEAGATE ST3300007FC
6	6 2(4)	14	17245	200MB	NONE	FRMT DRV	SEAGATE ST318453FC
7	2(4)	15	17304	200MB	NONE	NEW DRV	SEAGATE ST318304FC

6. Configure the parameters (see the subsequent sections for details).

7. A confirmation box will appear on the screen. Verify all information in the box before choosing Yes to confirm and proceed.

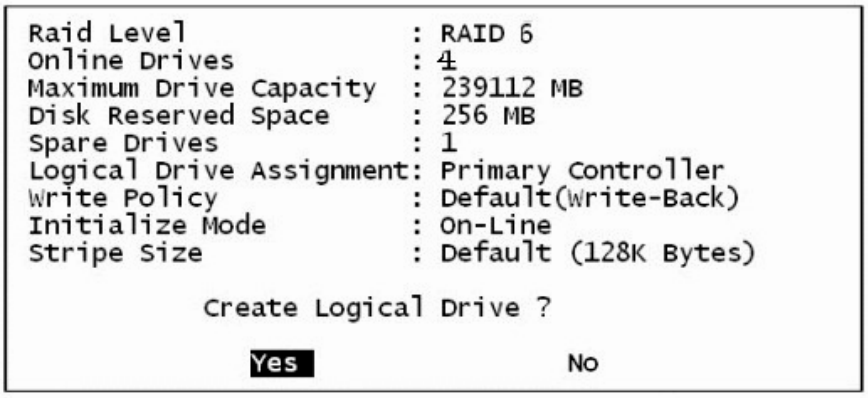
Raid Level : RAID 6
Online Drives : 4
Maximum Drive Capacity : 239112 MB
Disk Reserved Space : 256 MB
Spare Drives : 1
Logical Drive Assignment: Primary Controller
Write Policy : Default(Write-Back)
Initialize Mode : On-Line
Stripe Size : Default (128K Bytes)
Create Logical Drive ?
Yes
No

8. If the online initialization mode is applied, the logical drive will first be created and the controller will initialize the array in the background or when the array is less stressed by I/Os.

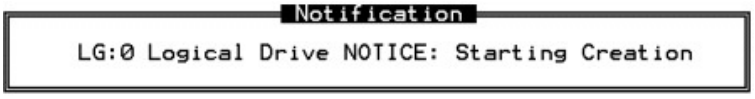

9. The completion of array creation is indicated by the message prompt above.

10. A controller event will prompt to indicate that the logical drive initialization has begun. Press ESC to cancel the “Notification” prompt, and a progress indicator will display on the screen as a percentage bar.

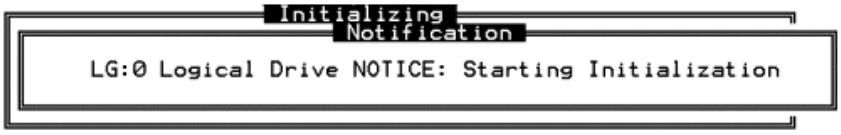

11. While the array initialization runs in the background, you can continue configuring your RAID subsystem, e.g., with host LUN mapping. When a fault-tolerant RAID level (RAID 1, 3, 5 or 6) is selected, the subsystem will start initializing parity.

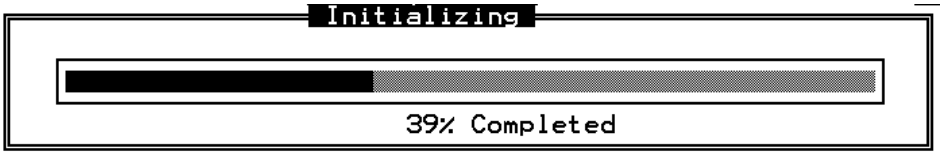

12. Use the ESC key to view the status of the created logical drive.

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	0	C	*LN	*SB	*FL	NAME
A0	C54758	0	RAID5	140000	GOOD			7	B	8	0	0	
A1	1E45E919	0	RAID5	120000	DRV FAILED			7	B	6	0	1	
B2	182648C2	1	RAID5	50000	GOOD			7	B	6	0	0	
B3	2AEB4D85	1	RAID5	50000	GOOD			7	B	6	0	0	
4													
5													
6													
7													

NOTE

Only logical drives with RAID levels 1, 3, 5, or 6 will take the time to initialize the logical drive. Logical drives with RAID level 0 and NRAID do not perform logical drive initialization. With RAID0 or NRAID, the drive initialization process finishes almost immediately.

Also, the Parity Regeneration function is also absent from a RAID0 or NRAID array

menu.

There will be a warning message if you want to create a logical drive larger than 64GB.

7.2.2 Setting Maximum Drive Capacity

1. After you selected the members, press [ESC] to proceed. A Logical Drive Preference menu will prompt.

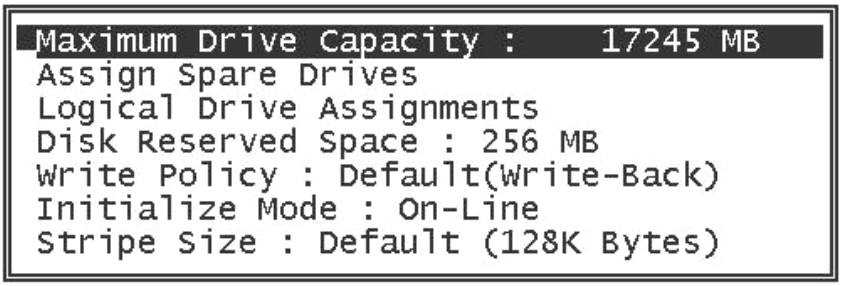
Maximum Drive Capacity : 17245 MB
Assign Spare Drives
Logical Drive Assignments
Disk Reserved Space : 256 MB
Write Policy : Default(Write-Back)
Initialize Mode : On-Line
Stripe Size : Default (128K Bytes)

2. As a rule, a logical drive should be composed of drives of the same capacity. A logical drive can only use the capacity of each drive up to the maximum capacity of the smallest member. The capacity of the smallest member will be listed here as the maximum drive capacity.

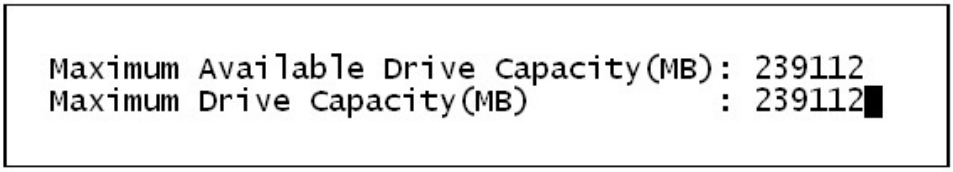

7.2.3 Assigning Spare Drives

NOTE

A logical drive composed in a non-redundancy RAID level (NRAID or RAID0) has no fault-tolerance and does not support spare drive rebuild.

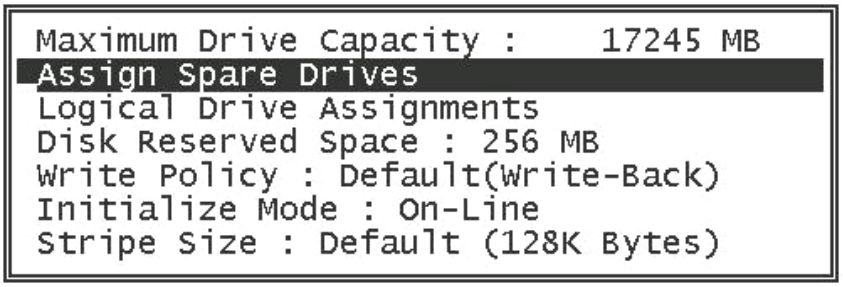
Maximum Drive Capacity : 17245 MB
Assign Spare Drives
Logical Drive Assignments
Disk Reserved Space : 256 MB
Write Policy : Default(Write-Back)
Initialize Mode : On-Line
Stripe Size : Default (128K Bytes)

1. You can assign a “Local Spare” drive to the logical drive from a list of unused disk drives. The spare chosen here is a spare exclusively assigned and will automatically replace a failed drive within the logical drive. The controller will then rebuild data onto the replacement drive in the event of a disk drive failure.

Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
2(4)	15	17304	200MB	NONE	NEW DRV	SEAGATE ST318304FC

2. The reserved space is a small section of disk space formatted for storing array configuration, Embedded RAIDWatch program, and other non-volatile data. This item is for display only - you cannot change the size of the reserved space.

NOTE

Assigning a hot-spare to an array composed of drives of a different interface type should be avoided. For example, a SATA Global spare may accidentally participate in the rebuild of an array using SAS members. It is better to prevent mixing SAS and SATA drives in a logical drive configuration because they have different rotation speeds and capacities.

7.2.4 Changing Logical Drive Assignments

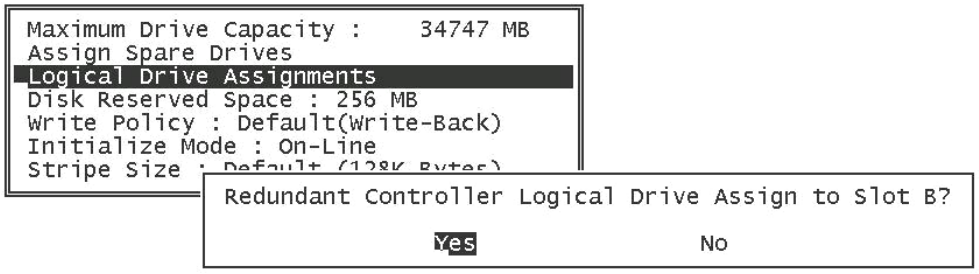
Maximum Drive Capacity : 34747 MB
Assign Spare Drives
Logical Drive Assignments
Disk Reserved Space : 256 MB
Write Policy : Default(Write-Back)
Initialize Mode : On-Line
Redundant Controller Logical Drive Assign to Slot B?
Yes
No

You do not need to change LD assignment here. Once logical drives are included into logical volumes, assign logical volumes to different controllers to balance workload.

7.2.5 Changing Write Policy

This sub-menu allows you to select the caching mode for this specific logical drive. “Default” is a neutral value that is coordinated with the subsystem’s general caching mode setting bracketed in the Write Policy status.

NOTE

The “Write-back” and “Write-through” parameters are permanent for specific logical drives. The “Default” selection, however, is more complicated and more likely equal to “not specified.”

If set to “Default,” a logical drive’s write policy is determined not only by the system’s general caching mode setting, but also by the “Event trigger” mechanisms. The “Event Trigger” mechanisms automatically disable the write-back caching and change to the conservative “Write-through” mode in the event of component failures or elevated temperature.

7.2.6 Setting the Initialization Mode

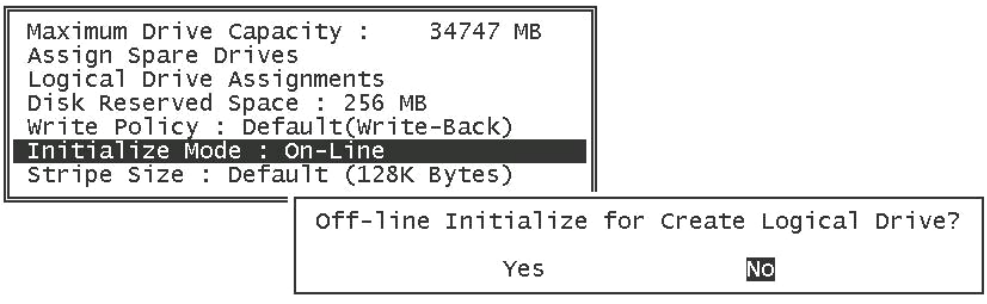
Maximum Drive Capacity : 34747 MB
Assign Spare Drives
Logical Drive Assignments
Disk Reserved Space : 256 MB
Write Policy : Default(Write-Back)
Initialize Mode : On-Line
Stripe Size : Default (128K Bytes)
Off-line Initialize for Create Logical Drive?
Yes
No

This sub-menu allows you to configure if the logical drive is immediately available. If the online (default) mode is used, logical drive is immediately ready for I/Os and you may continue with array configuration, e.g., partitioning the array, before the array’s initialization process is completed.

7.2.7 Setting the Stripe Size

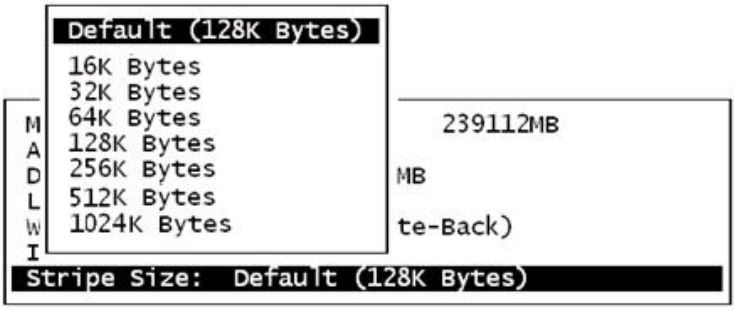

This option should only be changed by experienced technicians. Setting to an incongruous value can severely drag performance; therefore, this option should only be changed when you can be sure of the performance gains it might bring you.

For example, if your array is often stressed by large and sequential I/Os, a small stripe size will force hard disks to spin many more times in order to conduct data in different data blocks and hence reduce the efficiency brought by parallel executions. Diagrammed below are conditions featuring host I/Os in 512KB transfer size and a RAID3 array using 128KB and 32KB stripe sizes. The first condition shows a perfect fit where each host I/O is efficiently satisfied by writing to 4 disks simultaneously.

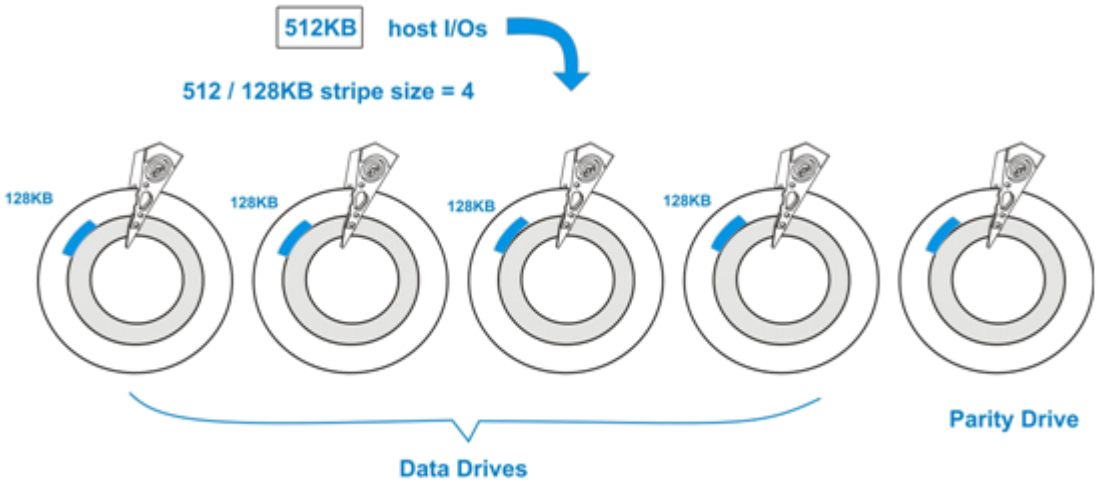

RAID3 Array

As the contrast, an inadequately small, 32KB stripe size will force the hard disks to write four times and controller firmware to generate 4 parity blocks.

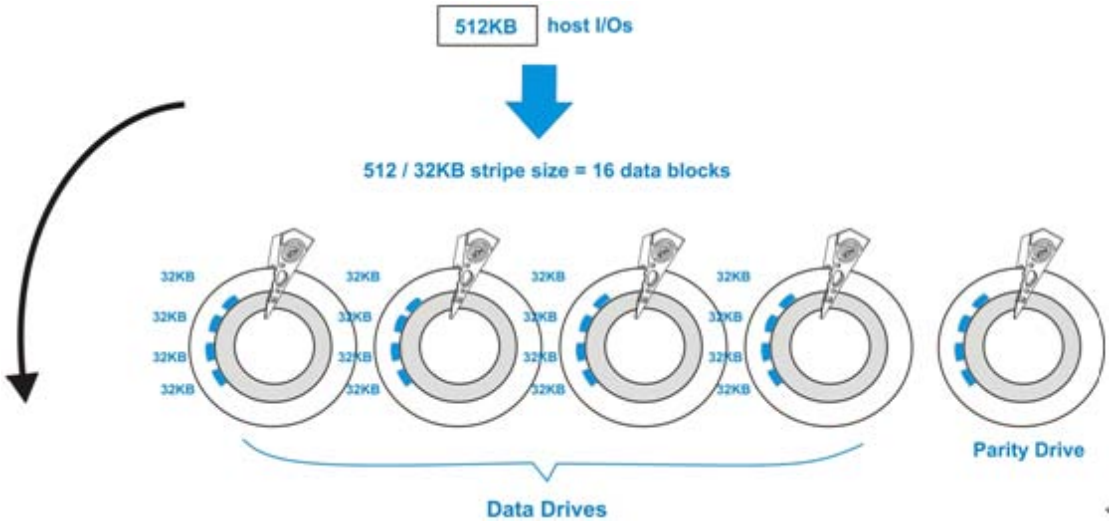

The “Default” value is determined by the combined factors of the controller Optimization Mode setting and the RAID level selected for the specific logical drive. See the table below for default values:

RAID Level	Stripe Size
RAID0	128KB
RAID1	128KB
RAID3	16KB
RAID5	128KB
RAID6	128KB
NRAID	128KB

Press [ESC] to continue when all the preferences have been set.

NOTE

The Stripe size here refers to the “Inner Stripe Size” specifying the chunk size allocated on each individual data disk for parallel access instead of the “Outer Stripe Size” that is the sum of chunks on all data drives.

7.2.8 Setting the Power Saving Mode

Go to: View and Edit Logical Drives > (Logical Drive) > Power Saving

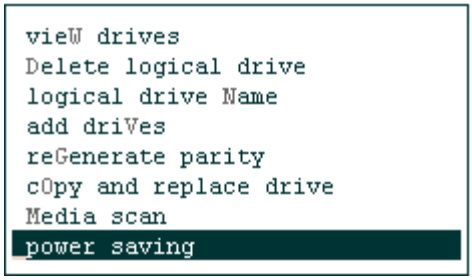

This feature supplements the disk spin-down function, and supports power-saving on specific logical drives or non-member disks such as spare drives. With no host I/Os, disk drives can consequentially enter two power-saving modes: Level 1 in idle mode and Level 2 in spun-down mode.

Applicable Disk Drives:

Logical drives and non-member disks [including spare drives and unused drives (new or formatted drives)]. The power-saving policy set to an individual logical drive (from the View and Edit Logical Drive menu) has priority over the general Drive-side Parameter setting.

Power-saving Levels and Features:

Level	Power Saving Ratio	Recovery Time	ATA command	SCSI command
Level 1	15% to 20%	1 second	Idle	Idle
Level 2	80%	30 to 45 seconds	Standby	Stop

NOTE

HDD vendors have different implementations for the idle mode. Most vendors ramp-load or park the hard disk actuator arm, while not all vendors reduce the rotation speed.

- Hard drives can be configured to enter the Level 1 idle state for a configurable period of time and then enter the Level 2 spin-down state.
- The combinations of power-saving modes can be:
  Disable,
  Level 1 only,
  Level 1 and then Level 2,
  Level 2 only. (Level 2 is equivalent to legacy spin-down)

- The Factory defaults is “Disabled” for all drives. The default for logical drives is also Disabled.
- The preset waiting period before entering the power-saving state:
  Level 1: 5 minutes (5 minutes without I/O requests)
  Level 2: 10 minutes (10 minutes since the beginning of level 1)
- If a logical drive is physically relocated to another enclosure (drive roaming), all related power-saving feature is cancelled.

7.2.9 Editing Logical Drives

Go to: View and Edit Logical Drives > (Logical Drive)

Select “View and Edit Logical Drives“ in the Main Menu to display the array status. Refer to the previous chapter for more details on the legends used in the Logical Drive’s status. To see the drive member information, choose the logical drive by pressing [ENTER].

The logical drive-related functions include:

View Drive	Displays member drive information
Delete Logical Drive	Deletes a logical drive.
Partition Logical Drive	Creates or removes one or more partition within a logical drive
Logical Drive Name	Assigns a name to a logical drive
Rebuild Logical Drive	Manually rebuilds a logical drive when a failed drive is replaced
Expand Logical Drive	Expands a logical drive using the unused capacity
Migrate Logical Drive	Migrates a logical drive to a different RAID level
Add Drives	Adds physical drive(s) to a logical drive
Regenerate Parity	Regenerates a logical drive’s parity
Copy and Replace Drive	Copies or replaces members of a logical drive
Media Scan	Configures Media Scan priority, iteration count, and task schedules

Write Policy	Changes the write policy associated with the logical drive
Power Saving	Activates power saving mode.

7.2.10 Deleting a Logical Drive

Go to: View and Edit Logical Drives > (Logical Drive) > Delete Logical Drive

NOTE

Deleting a logical drive destroys all data stored on it.

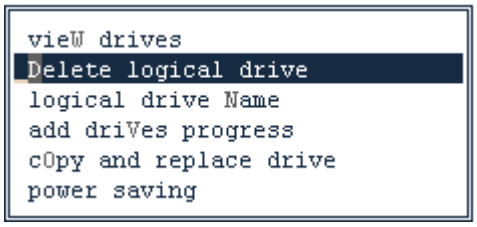

7.2.11 Naming a Logical Drive

Go to: View and Edit Logical Drives > (Logical Drive) > Logical Drive Name

Naming can help identify different arrays in a multi-array configuration.

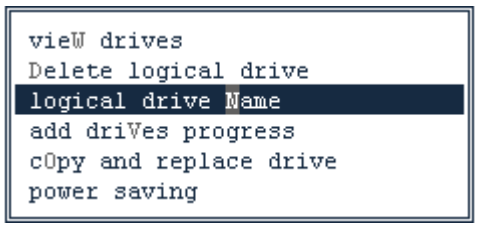

NOTE

This function is especially helpful in situations such as the following:

One or more logical drives have been deleted, the array indexing is changed after system reboot, e.g., LD0 deleted and the succeeding LD1 becomes LD0. The designating numbers of logical drives following a deleted configuration will all be affected.

The maximum number of characters for a logical drive name is 32.

7.2.12 Expanding a Logical Drive or a Logical Volume

Go to: View and Edit Logical Drives > (Logical Drive) > Expand Logical Drive

A logical volume can only be expanded if its subordinate logical drives have free, unused capacity.

There are several conditions for logical drives to have free capacity:

- The Max. Drive Capacity for each member drive has intentionally been reduced while creating a logical drive. Some capacity can be hidden and left intentionally unused by setting the Max. Drive Capacity to a lower number.
- Logical drive capacity has been expanded by adding new drives. To do that, you must have free drive bays in your enclosure.
- Logical drive capacity has been expanded by copying and replacing drive members. Using drives of larger capacity, you can replace the original members of a logical drive.

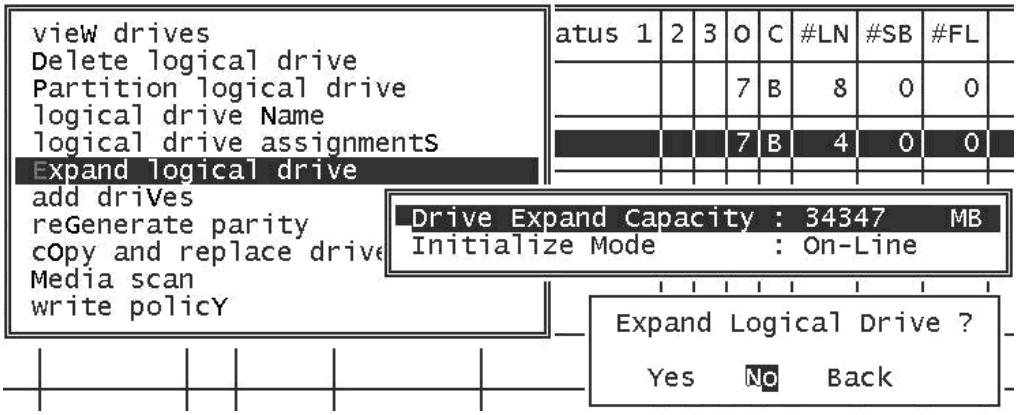

NOTE

The Drive Expand Capacity here refers to the unused capacity on each member drive. If a RAID5 array has 4 members and each member drive features a 2GB unused capacity, then the total unused capacity will be 4 - 1 (parity drive) x 2G = 6GB.

The capacity brought by the array expansion process will be available as a “new” partition.

Once all logical drives within a logical volume are expanded, you can expand a logical volume. It is preferred that all logical drives included in a logical volume are configured using the same capacity drives and number of drives. Since logical drives

are striped together to form a larger logical volume, every member must have an identical free capacity in order to expand a logical volume. Striping requires the “least common denominator” approach to combine logical drives into a logical volume.

7.3 RAID Migration

Currently the RAID migration function supports the migration between RAID5 and RAID6.

Before proceeding with RAID migration, make sure you have sufficient free capacity or unused disk drives in your RAID array. RAID6 arrays require at least four (4) member drives and use additional capacity for the distribution of secondary parity. For example, if you want to migrate a RAID5 array consisting of three (3) drives to RAID6, one additional disk drive should be available.

Different features of RAID5 and RAID6 arrays are summarized as follows:

<table>
<tr><th></th><th>RAID5</th><th>RAID6</th></tr>
<tr><td>Min. No. of Member Drives</td><td>3</td><td>4</td></tr>
<tr><td rowspan="4">Usable Capacity</td><td>N-1 (1 drive’s capacity used for storing parity data)</td><td>N-2 (2 drives’ capacity used for storing parity data); N>=4</td></tr>
<tr><td colspan="2">If individual disk capacity = 100G,</td></tr>
<tr><td colspan="2">Capacity of a 4-drive RAID5 = (4 -1) x 100G = 300G</td></tr>
<tr><td colspan="2">Capacity of a 4-drive RAID6 = (4 -2) x 100G = 200G</td></tr>
<tr><td>Redundancy</td><td>Single disk drive failure</td><td>2 disk drives to fail at the same time</td></tr>
</table>

7.3.1 Requirements for Migrating a RAID5 Array

The precondition for migrating a RAID5 array to RAID6 is:

The “usable capacity” (instead of the sum of raw capacity) of the RAID6 array should be equal or larger than the “usable capacity” of the original RAID5 array.

To obtain a larger capacity for migrating to RAID6, you can:

- Add Drive(s): Include one or more disk drives into the array.

- Copy and Replace: Use larger disk drives in the array to replace the original members of the RAID5 array.

7.3.2 Migration Methods

The conditions for migrating a RAID5 array to RAID6 are diagrammed as follows:

Fault condition:

The usable capacity of the to-be RAID6 array is smaller than the usable capacity of the original RAID5 array.

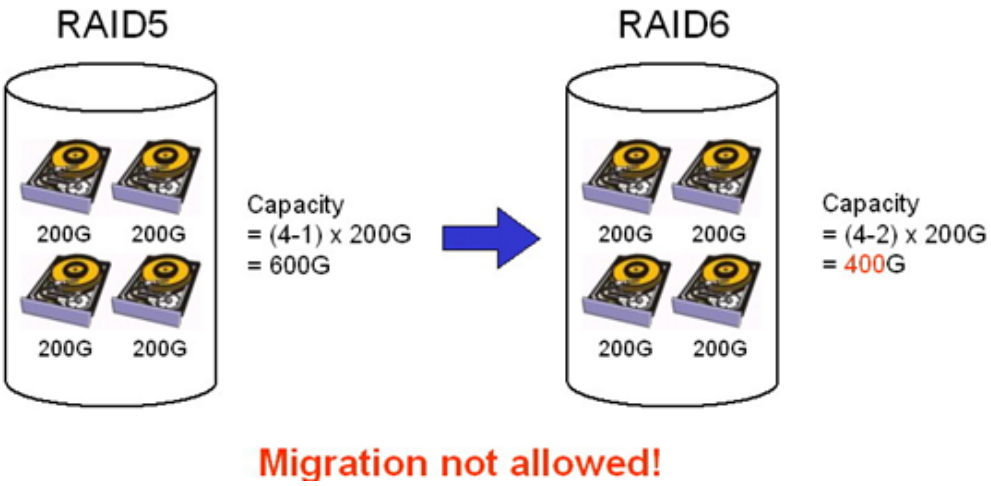

Migration by Adding Drive(s):

The additional capacity for migrating to a RAID6 array is acquired by adding a new member drive.

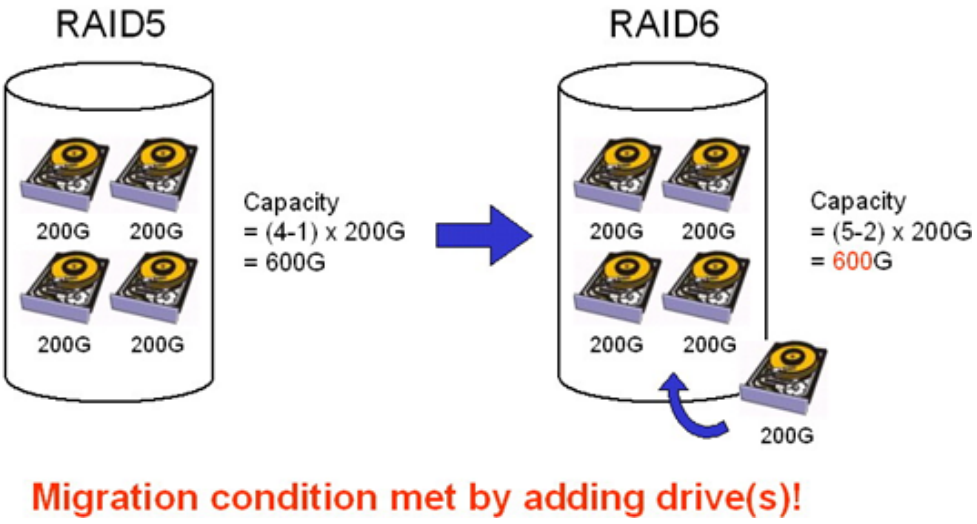

Migration by Copy and Replace:

The additional capacity for composing a RAID6 array is acquired by using larger

drives as the members of the array. Members of an existing logical drive can be manually copied and replaced using the “Copy & Replace” function.

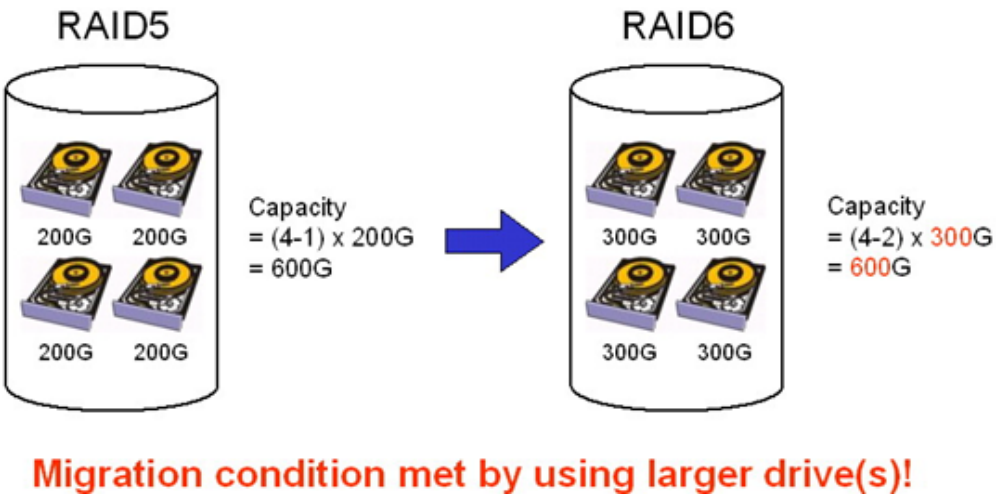

7.3.3 Migration: Exemplary Procedure

Go to: View and Edit Logical Drives > (Logical Drive) > Migrate Logical Drive

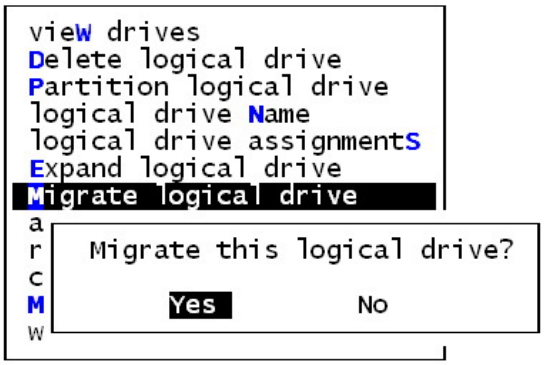

1. A selection box should prompt allowing you to choose a RAID level to migrate to. Press [ENTER] on RAID6.

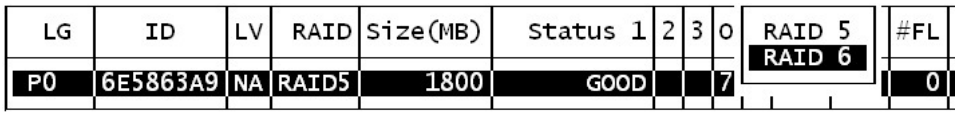

2. A list of member drives and unused disk drives (new or used drives) should prompt. In the case of migrating a 3-drive RAID5 to 4-drive RAID6, you can select the original members of the RAID5 array and select one more disk drive to meet the minimum requirements of RAID6. You may also select unused drives in your enclosure for composing the new RAID6 array.

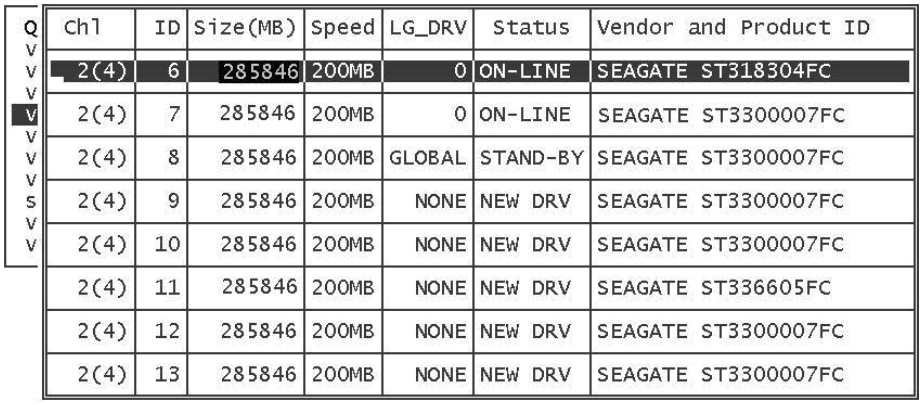

Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
2(4)	6	285846	200MB	0	ON-LINE	SEAGATE ST318304FC
2(4)	7	285846	200MB	0	ON-LINE	SEAGATE ST3300007FC
2(4)	8	285846	200MB	GLOBAL	STAND-BY	SEAGATE ST3300007FC
2(4)	9	285846	200MB	NONE	NEW DRV	SEAGATE ST3300007FC
2(4)	10	285846	200MB	NONE	NEW DRV	SEAGATE ST3300007FC
2(4)	11	285846	200MB	NONE	NEW DRV	SEAGATE ST336605FC
2(4)	12	285846	200MB	NONE	NEW DRV	SEAGATE ST3300007FC
2(4)	13	285846	200MB	NONE	NEW DRV	SEAGATE ST3300007FC

3. Press [ESC] to proceed to the next configuration screen. A sub-menu should prompt.

4. You may either change the maximum capacity to be included in the new RAID6 array or change the array stripe size.

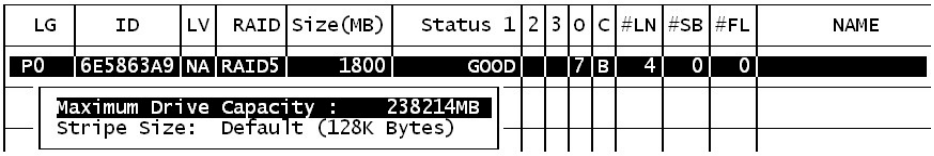
LG ID LV RAID Size(MB) Status 1 2 3 0 C #LN #SB #FL NAME
P0 6E5863A9 NA RAID5 1800 GOOD 7 B 4 0 0
Maximum Drive Capacity : 238214MB
Stripe Size: Default (128K Bytes)

5. A confirmation box should prompt. Check the configuration details and select Yes to start the migration process.

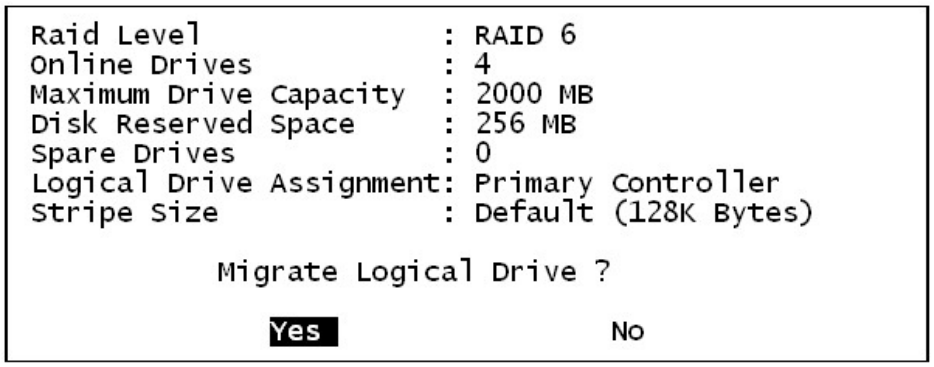
Raid Level : RAID 6
Online Drives : 4
Maximum Drive Capacity : 2000 MB
Disk Reserved Space : 256 MB
Spare Drives : 0
Logical Drive Assignment: Primary Controller
Stripe Size : Default (128K Bytes)
Migrate Logical Drive ?
Yes
No

6. A message should prompt indicating the migration process has started.

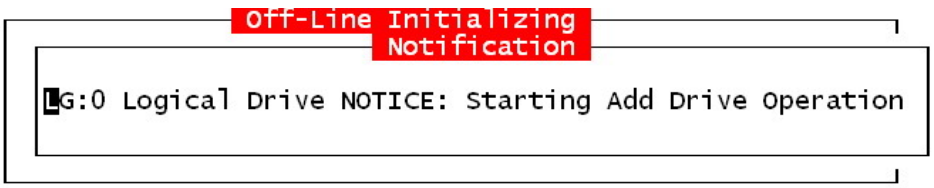

7. Press [ESC] to clear the message. The initialization progress is shown below.

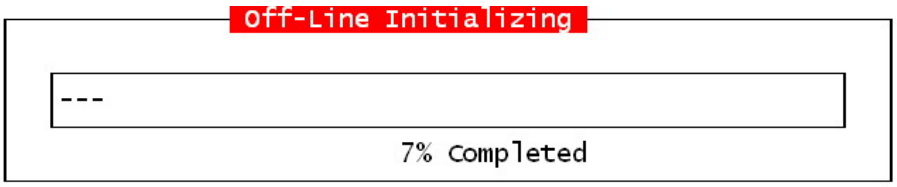

8. Since the migration process includes adding a new member drive, the completion of RAID migration is indicated as follows:

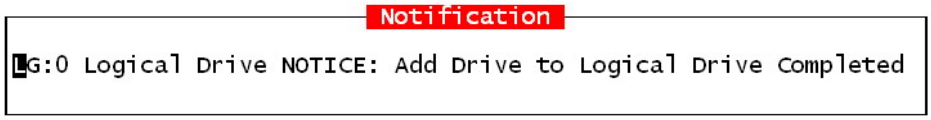

9. Once the migration is completed, associate the RAID6 array with the ID/LUN number originally associated with the previous RAID5 array.

7.4 Working with Logical Volumes

7.4.1 Creating a Logical Volume (Required)

Go to: View and Edit Logical Volumes

LV	ID	RAID	Size(MB)	Status 1	2	3	C	#LN	#FL	Name
A0	06FBCA387F22D6EF	RAID0	238214	GOOD			B	1	0	LV0
1										
2										
3										
4										
5										
6										
7										

NOTE

Unlike older Galaxy HDX, HDX2 and HDX3, you must create logical volumes to enwrap logical drives, and then create logical partitions as LUNs.

A logical volume consists of one or several logical drives. The member logical drives are striped together.

It is recommended to select logical drives identical in sizes and drive characteristics

into a logical volume, e.g., 2 RAID5 logical drives made of 8 SAS 10k rpm members.

1. Select “View and Edit Logical Volumes” in the Main Menu to display the current logical volume configuration and status on the screen. Select a logical volume index number (0 to 7) that has not yet been defined, and then press [ENTER] to proceed.

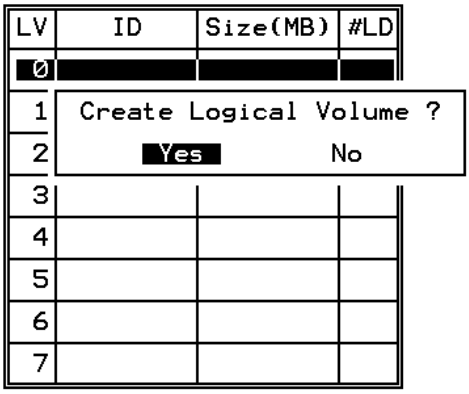

2. A prompt “Create Logical Volume?” will appear. Select Yes and press [ENTER].

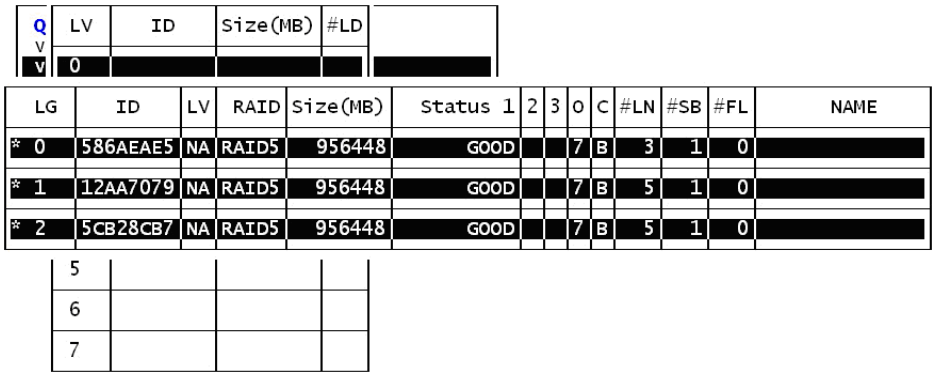

3. Select one or more logical drive(s) available on the list. The same as creating a logical drive, the logical drive(s) can be tagged for inclusion by positioning the cursor bar on the desired disk drive and pressing [ENTER] to select. An asterisk (*) will appear on the selected logical drive. Pressing [ENTER] again will deselect a logical drive.

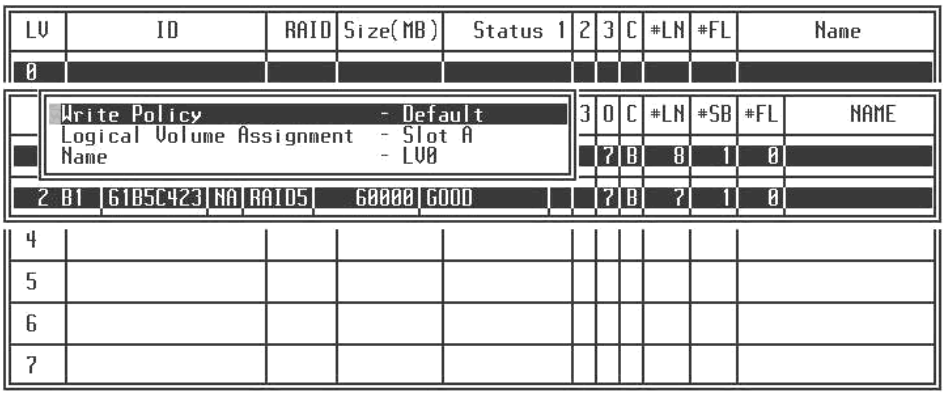

4. Use the arrow keys to select a sub-menu and change the write policy, controller assignment, and the name for the logical volume. You can balance the workload on partner RAID controllers by assigning volumes to both of them, e.g., 2 to Slot A controller and 2 volumes to the Slot B controller.

5. Logical volumes can be assigned to different controllers (primary or secondary; Slot A or Slot B controllers). The default is the primary or Slot A controller.

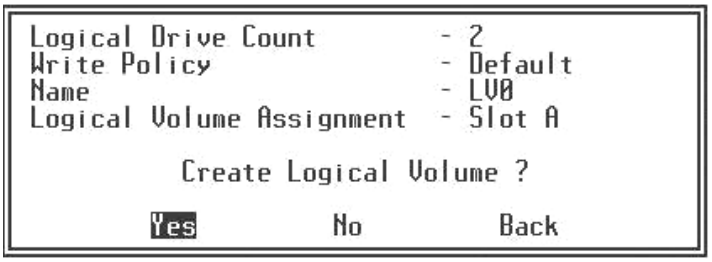

Note that if a logical volume is manually assigned to a specific controller, all its members' assignments will also be shifted to that controller.

6. When all the member logical drives have been selected, press [ESC] to continue. Choose Yes to create the logical volume.

7. Press [ENTER] on a configured volume, and the information of the created logical volume displays.

LV:	Logical Volume ID
ID:	Unique ID for the logical volume, randomly generated by the RAID controller firmware
RAID:	Members are striped together, always shows RAID0
Size:	Capacity of this volume
#LN:	Number of the included members
LV:	Logical Volume ID

7.4.2 Notes on Partitions in Galaxy HDX4 Series

- For Galaxy HDX4, logical partitions are the basic LUN units. Unlike older Galaxy models, you can not create partitions in logical drives. You should include logical drives into logical volumes, create logical partitions from volumes, and map

partitions to host as LUNs.

- If operating in a Unix-based system, reset the subsystem for the configuration changes to take effect if any changes were made to partition sizes and partition arrangement.

Below are the differences between older Galaxy partitions & Galaxy HDX4 partitions:

- Older Galaxy: Partitioning is used for slicing usable capacity in a volume. Changes in partition sizes affect other partitions.
- Galaxy HDX4: Partitioning is like creating LUN units from a storage pool. The unused capacity in pool can be reserved for keeping snapshot images, metadata, and configuration data for volume copy and volume mirror.
- Volume Copy and Volume Mirror take place between logical partitions. These functions are exerted on the Galaxy Array Manager GUI.
- You cannot expand a logical volume by adding new logical drives in the HDX4 series.
- In HDX4, a logical volume can be expanded by adding disk drives to its member logical drives or by replacing drives in its member logical drives.

7.4.3 Creating a Partition from Logical Volume (Required)

Go to: View and Edit Logical Volumes > Partition Logical Volume

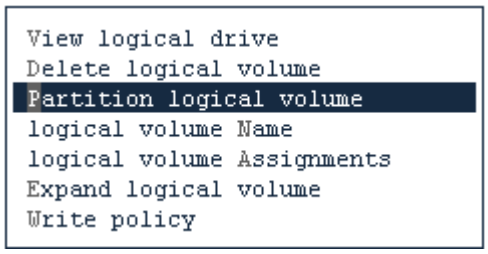

1. Press [ENTER]. Select Yes and press [ENTER] to proceed.

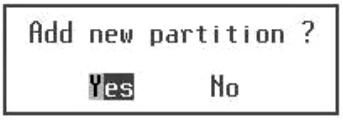

2. A Capacity window will prompt. Press [ENTER] on it to change the capacity for the partition.

3. Key in the desired capacity for the selected partition, and then press [ESC] to proceed. The remaining capacity will be a free space which can be leveraged for creating more partitions or keeping snapshots.

4. Move cursor bar to the Partition Name and change the partition name. Changing names is highly recommended especially for complex storage configurations.

5. You will prompted again by a confirm box. Choose Yes to confirm then press [ENTER].

6. The below message will prompt and the creation may take several seconds.

7. A list of partitions will appear. Press [ENTER] on any of the existing partitions to bring out a command menu where you can Add new Partition, change partition name, expand partition, or Delete a partition.

8. Follow the same procedure to create more partitions from your logical volumes as long as they have free capacity. Up to 64 partitions can be created from a single volume.

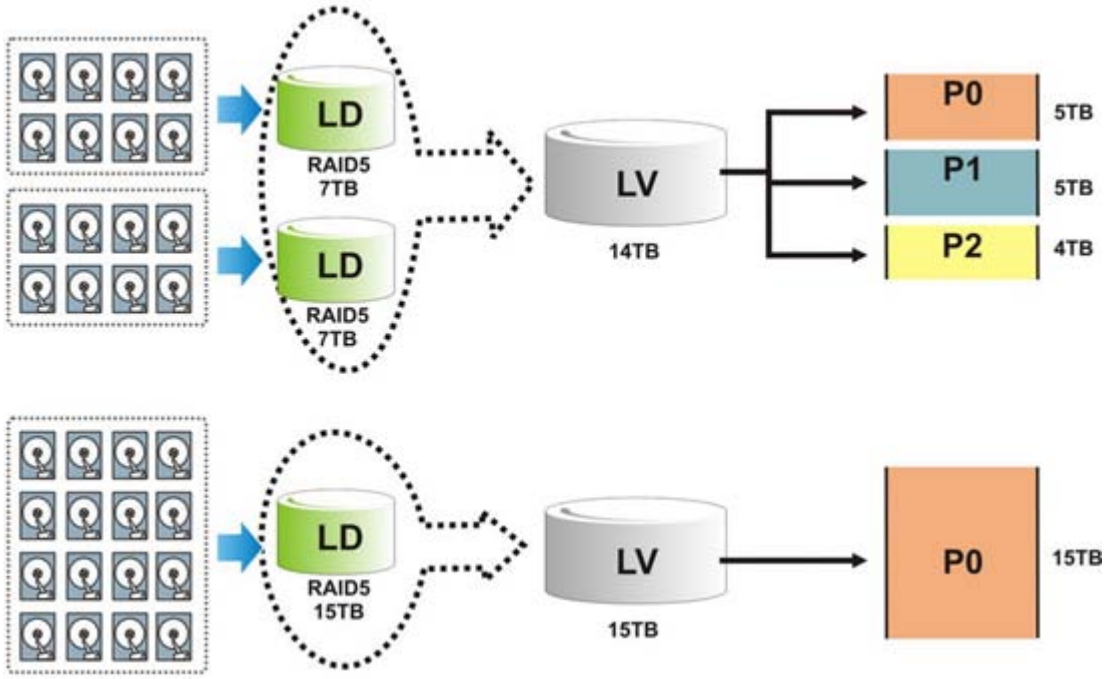

If you plan to protect your data using the snapshot function, make sure you leave enough space within a volume for keeping the snapshot images. Depending on the sizes and frequency of data changes in a source volume, snapshots can take up enormous disk space.

7.4.4 Deleting Partitions

Go to: View and Edit Logical Volumes > Partition Logical Volume > Delete Partition

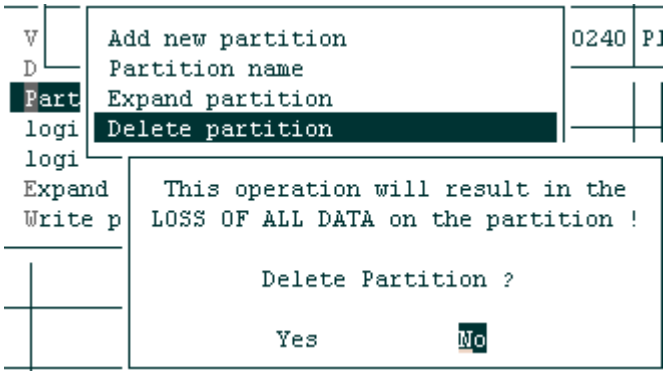

Unlike the standard older Galaxy HDX, HDX2, and HDX3, the disk space of a deleted partition will be automatically allocated to the free space partition as diagrammed below. For example, if partition 1 is deleted, its disk space will be released to the free space.

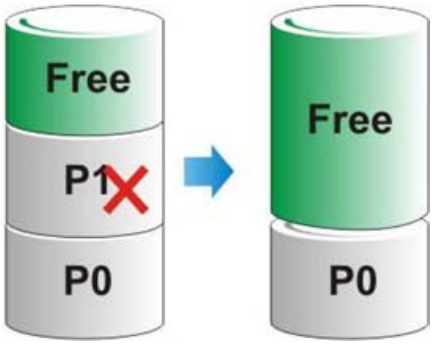

7.4.5 Examining Valid Connectivity

Go to: View and Edit Channels > View Device Port Name List

Before you begin LUN mapping, you should have host links properly connected, and if necessary, configured VLAN (using iSCSI storage) or switch zoning.

Normally, if a Fibre Channel HBA port has an access route to controller host ports, it should be listed on the port name list. Shown below is an example of a list of FC HBA ports.

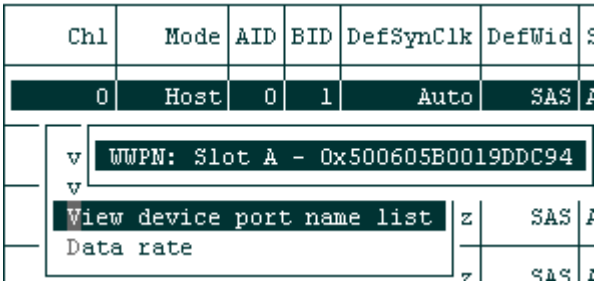

If HBA ports do not appear on the list, configuration or cabling faults could have occurred in your environment.

You should know the port names of HBA ports on your application servers via the HBA BIOS utility, a SAN exploring software, or check on the HBA port name label.

If properly configured, HBA port names should all appear on the port name list.

Below is a basic deployment. You should sketch your topology with specified links and details like FC port names.

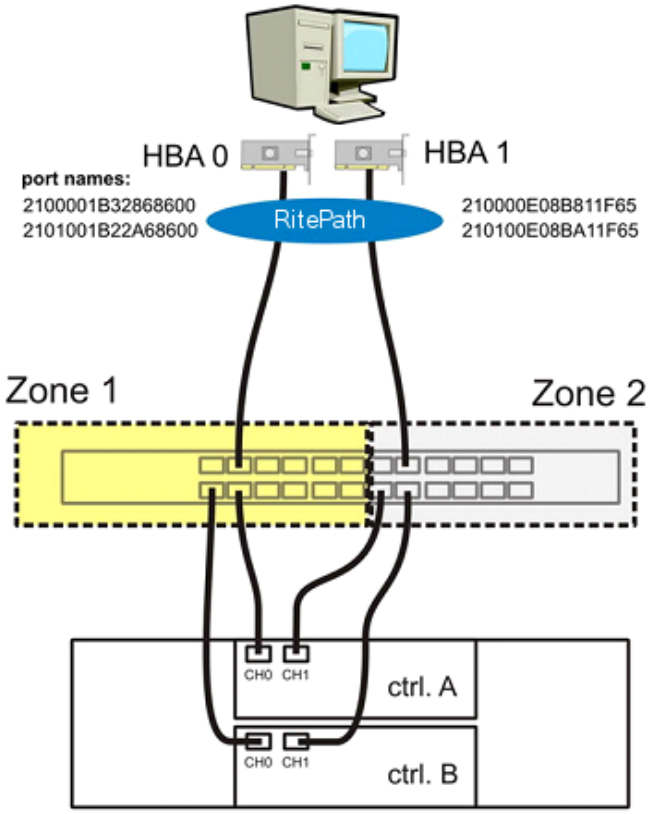

7.4.6 Managing Host Adapter Ports

Go to: View and Edit Channels > View Device Port Name List > (Port)

You can assign a nickname to a HBA port. This eases identification of multiple servers and its FC ports in a complex configuration.

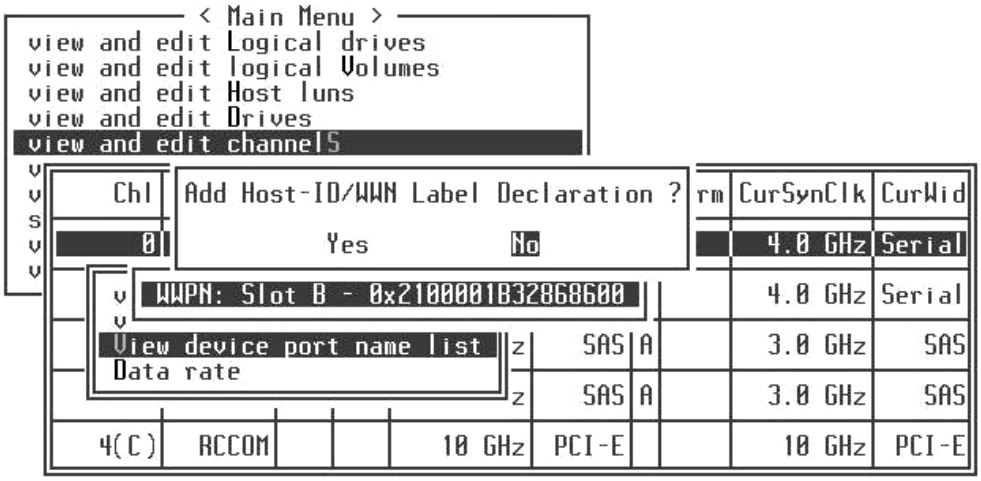

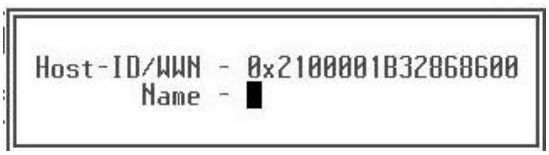

Select an HBA port name and key in a name for a specific FC port, e.g., server1port1. Use only numeric and alphabetic characters.

The same list maintenance options can also be found in “view and edit host luns” -> “Edit Host-ID/WWN Name List.”

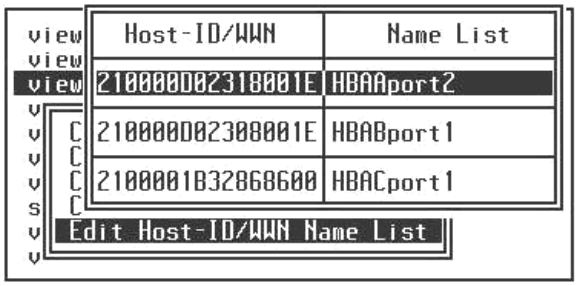

You can create a list of HBA port names by polling detected names or manually key in names on servers that are currently not powered-on.

You should also develop a drawing of logical associations showing how your storage volumes are tied with host channel ID/LUNs.

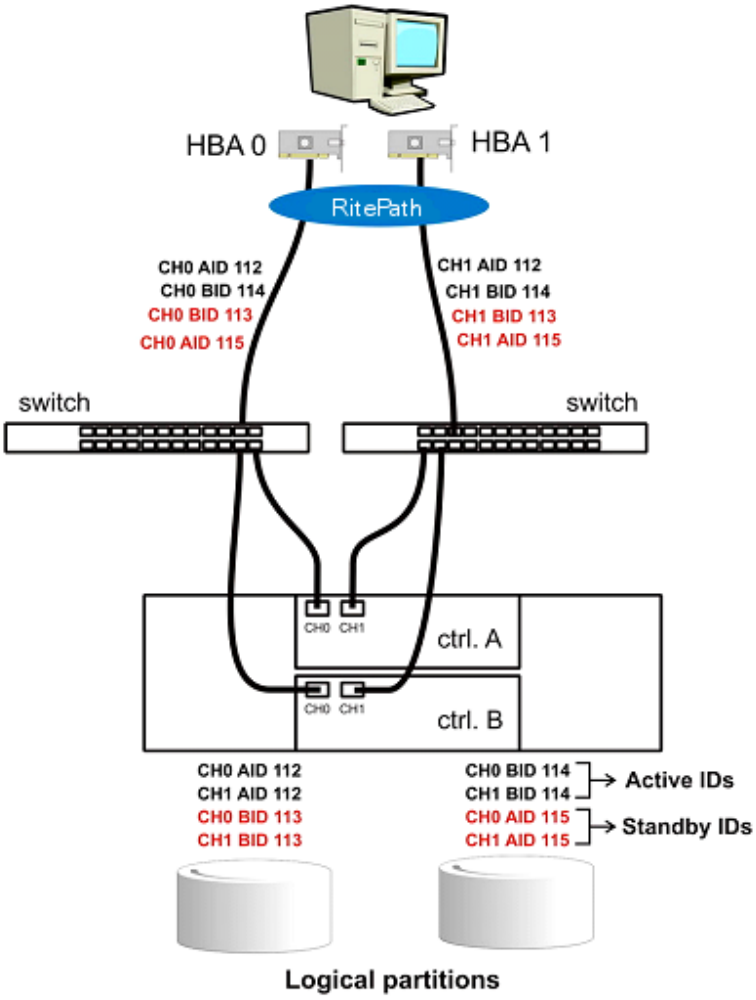

7.4.7 Notes on Mapping a Partition to Host LUN

- Your subsystem comes with 1 Slot A ID and 1 Slot B ID only. You need to manually create more host channel IDs in a dual-controller configuration. Please enter “View and Edit Channels” menu to create or remove a host ID.
- The latest firmware revision supports the cross-controller ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs. However, mapping to both controllers’ IDs is usually beneficial when it is difficult making the fault-tolerant host links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. Cross-controller mapping also makes sense in clustered-server environments. Currently, no external SAS switch has not gained popularity on the market. For Fibre-host systems, fault-tolerant links can easily be made with the help of external bypass such as Fibre Channel switches.
- If your host adapter cards do not support multiple LUN numbers under a channel

ID, select LUN0. You should refer to the documentation that came with your host adapters to see whether multiple LUNs are available.

The differences between Map Host LUN and Extended LUN:

- Map Host LUN simply presents a logical partition to the host links. If host links are made via an FC switch, all servers attached to the switch (or those within the same zone) can “see” the partition.
- The Extended LUN mapping binds a logical partition with a specific HBA port and presents the partition to the HBA port.

For examples of fault-tolerant link connections, please refer to your system Hardware Manual.

7.4.8 Mapping a Partition to a Host LUN

1. Select “View and Edit Host luns“ in the Main Menu, then press [ENTER]. A list of host channel IDs will appear.

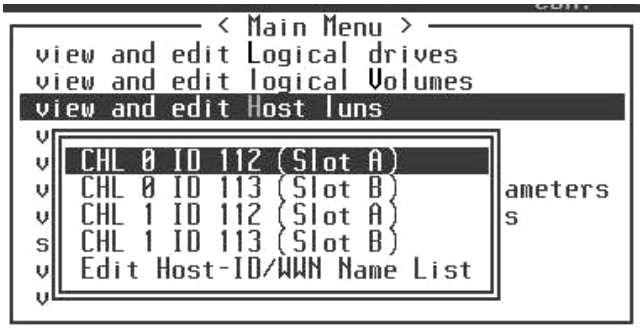
< Main Menu >
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
CHL 0 ID 112 (Slot A)
CHL 0 ID 113 (Slot B)
CHL 1 ID 112 (Slot A)
CHL 1 ID 113 (Slot B)
Edit Host-ID/WWN Name List
ameters
s

2. A list of host channel/ID combinations appears on the screen. The diagram above shows two host channels and each is designated with at least a default ID. More can be manually added on each channel.

3. Multiple IDs on host channels are necessary for creating access to RAID arrays through fault-tolerant data links. Details on creating multiple IDs and changing channel modes have been discussed in the previous chapter. Select a host ID by pressing [ENTER].

4. Select the channel-ID combination you wish to map, then press [ENTER] to proceed. An index of LUN numbers will display. Select an LUN number under the ID. Press [ENTER] on an LUN number to proceed and press [ENTER] again on “Map Host LUN” or “Extended LUN Mapping.”

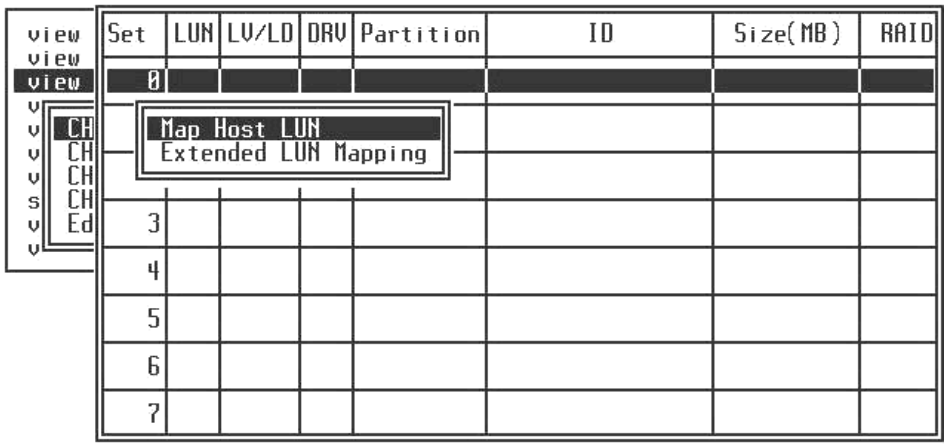

5. All existing logical volumes will be listed. Select a logical volume by pressing [ENTER] on it.

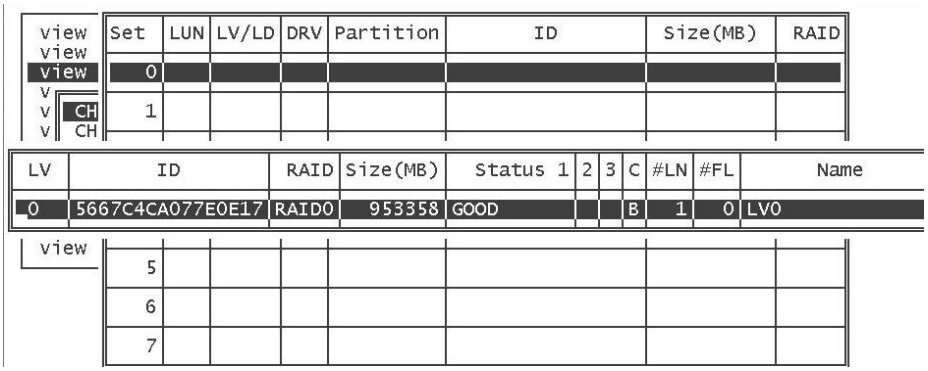

6. Logical partitions within that volume will be listed. Select a partition by moving the cursor bar and press [ENTER] on it.

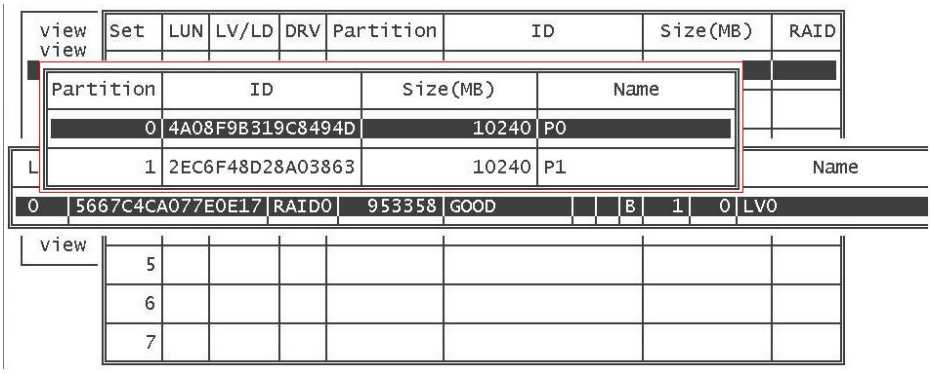

7. When prompted by the confirmation message, check the mapping details and select Yes to complete the process.

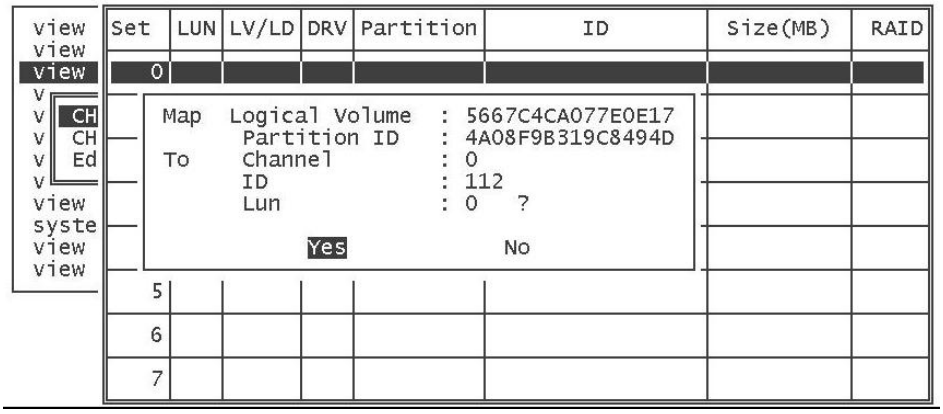

8. The details in the confirmation box read: partition “xxxxxxxx8494D” of logical volume “xxxxxxE0E17” will map to (be associated with) LUN 0 of ID 112 on host channel 0.

9. Repeat the process to complete host LUN mapping.

> NOTE
>
> Once any host ID/LUN is successfully associated with a logical partition, the “No Host LUN” message in the LCD screen will change to “Ready.”

7.4.9 Deleting Host LUNs

Go to: View and Edit Host Luns > (Channel) > (LUN)

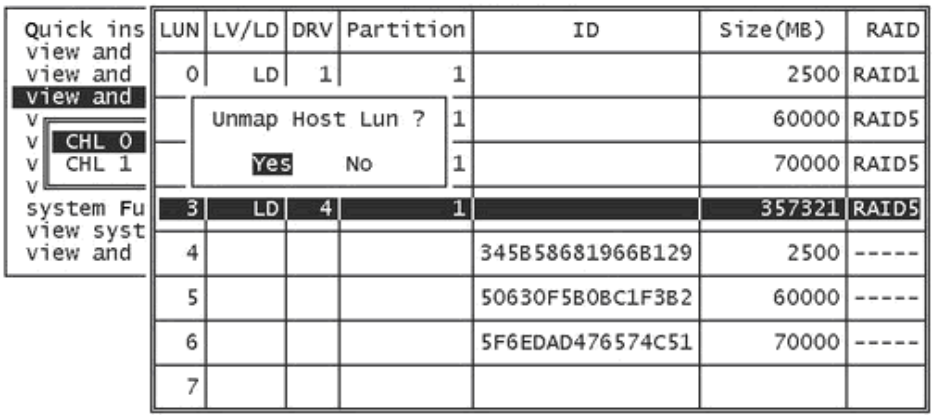

7.4.10 Expanding a Logical Volume

Go to: View and Edit Logical Drives > Expand Logical Drive

A logical volume can only be expanded if its subordinate logical drives have free, unused capacity.

There are several conditions for logical drives to have free capacity:

- The Max. Drive Capacity for each member drive has intentionally been reduced while creating a logical drive. Some capacity can be hidden and left intentionally unused by setting the Max. Drive Capacity to a lower number.
- Logical drive capacity has been expanded by adding new drives. To do that, you must have free drive bays in your enclosure.
- Logical drive capacity has been expanded by copying and replacing drive members. Using drives of larger capacity, you can replace the original members of a logical drive.

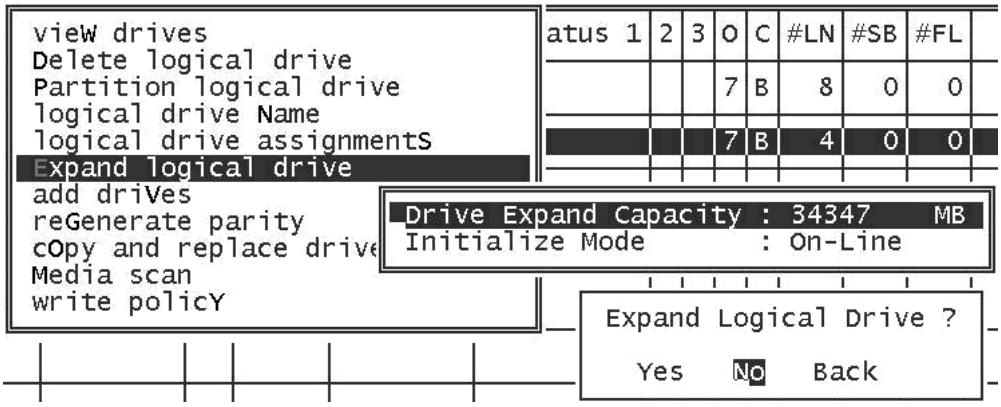

NOTE

The Drive Expand Capacity here refers to the unused capacity on each member drive. If a RAID5 array has 4 members and each member drive features a 2GB unused capacity, then the total unused capacity will be 4 - 1 (parity drive) x 2G = 6GB.

The capacity brought by the array expansion process will be available as a “new” partition.

Once all logical drives within a logical volume are expanded, you can expand a logical volume. It is preferred that all logical drives included in a logical volume are configured using the same capacity drives and number of drives. Since logical drives are striped together to form a larger logical volume, every member must have an identical free capacity in order to expand a logical volume. Striping requires the “least common denominator” approach to combine logical drives into a logical volume.

7.5 Assigning Spare Drive and Rebuild Settings

7.5.1 Adding a Local Spare Drive

Go to: View and Edit Drives > (Unassigned Drive) > Add Local Spare Drive

A spare drive is a standby drive that automatically participates in the rebuild of logical arrays. A spare drive must have an equal or larger capacity than the array members. A Local Spare is one that participates in the rebuild of a logical drive it is assigned to. A Global Spare participates in the rebuild of all configured logical drives, and it should have a capacity equal to or larger than all member drives in a RAID subsystem.

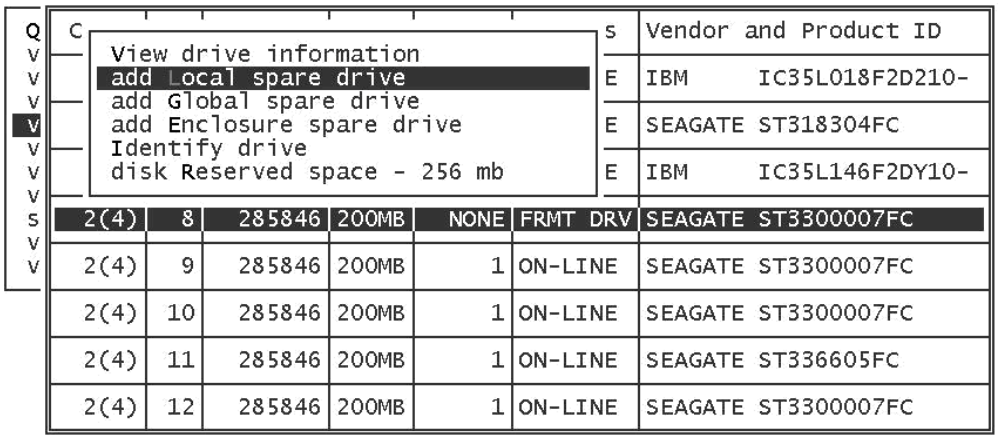
View drive information
add Local spare drive
add Global spare drive
add Enclosure spare drive
Identify drive
disk Reserved space - 256 mb
Vendor and Product ID
IBM IC35L018F2D210-
SEAGATE ST318304FC
IBM IC35L146F2DY10-
2(4) 8 285846 200MB NONE FRMT DRV SEAGATE ST3300007FC
2(4) 9 285846 200MB 1 ON-LINE SEAGATE ST3300007FC
2(4) 10 285846 200MB 1 ON-LINE SEAGATE ST3300007FC
2(4) 11 285846 200MB 1 ON-LINE SEAGATE ST336605FC
2(4) 12 285846 200MB 1 ON-LINE SEAGATE ST3300007FC

7.5.2 Adding a Global Spare Drive

Go to: View and Edit Drives > (Unassigned Drive) > Add Global Spare Drive

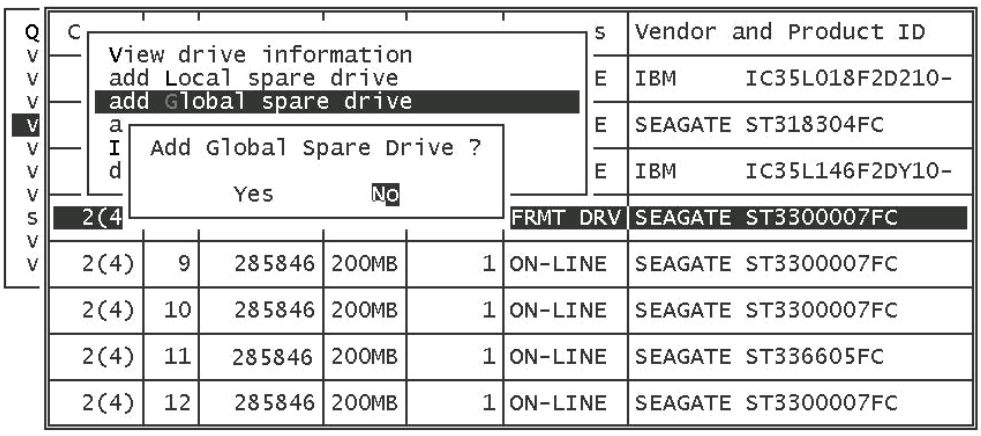
View drive information
add Local spare drive
add Global spare drive
Add Global Spare Drive ?
Yes No
Vendor and Product ID
IBM IC35L018F2D210-
SEAGATE ST318304FC
IBM IC35L146F2DY10-
FRMT DRV SEAGATE ST3300007FC
2(4) 9 285846 200MB 1 ON-LINE SEAGATE ST3300007FC
2(4) 10 285846 200MB 1 ON-LINE SEAGATE ST3300007FC
2(4) 11 285846 200MB 1 ON-LINE SEAGATE ST336605FC
2(4) 12 285846 200MB 1 ON-LINE SEAGATE ST3300007FC

7.5.3 Adding an Enclosure Spare Drive

Go to: View and Edit Drives > (Unassigned Drive) > Add Enclosure Spare Drive

An Enclosure Spare only participates in the rebuild of a failed drive located within the same enclosure.

NOTE

An Enclosure Spare is one that is used to rebuild a failed drive that resides in the same enclosure.

In configurations that span across multiple enclosures, a Global Spare may participate in the rebuild of a failed drive in a different enclosure. If rebuild takes place on drives across different enclosures, logical drives will end up having members dispersed in different enclosures. This results in management problems. For example, system administrators can mistakenly replace the wrong drive if yet another failure occurs. Using Enclosure Spare can avoid disorderly locations of member drives in a multi-enclosure configuration.

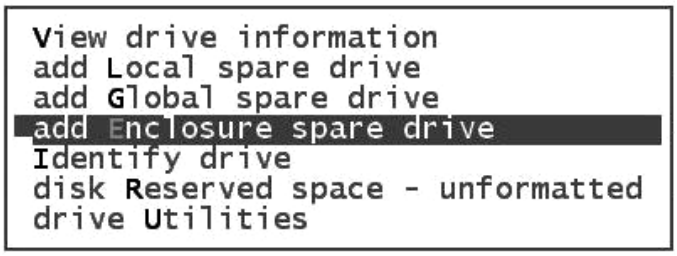

7.5.4 Deleting Spare Drive (Global/Local/Enclosure Spare Drive)

Go to: View and Edit Drives > (Spare Drive) > Delete Spare Drive

NOTE

The spare drive you deleted (disassociated or reassigned as a normal disk drive) or any drive you replaced from a logical unit will be indicated as a "used drive."

8 Fibre Channel Options

8.1 Viewing and Editing Channels

Go to: View and Edit Channels

The Galaxy HDX4 series come with preset data paths and there is no need to modify channel configurations, e.g., channel mode.

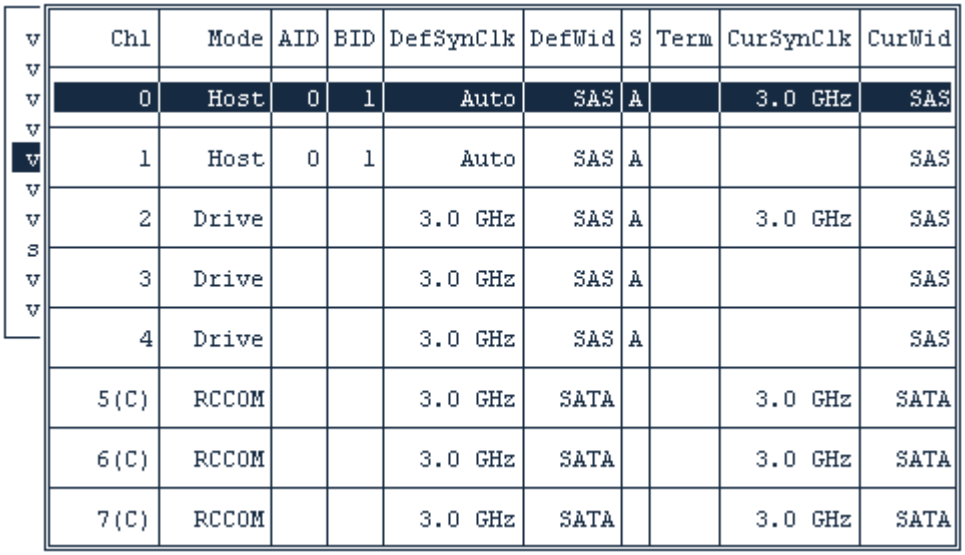

Chl	Mode	AID	BID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	0	1	Auto	SAS	A		3.0 GHz	SAS
1	Host	0	1	Auto	SAS	A			SAS
2	Drive			3.0 GHz	SAS	A		3.0 GHz	SAS
3	Drive			3.0 GHz	SAS	A			SAS
4	Drive			3.0 GHz	SAS	A			SAS
5(C)	RCCOM			3.0 GHz	SATA			3.0 GHz	SATA
6(C)	RCCOM			3.0 GHz	SATA			3.0 GHz	SATA
7(C)	RCCOM			3.0 GHz	SATA			3.0 GHz	SATA

NOTE

The Galaxy HDX4 models come with dedicated PCI-E channels that are strung between partner RAID controllers. These channels have no external interfaces and cannot be repurposed for I/Os. Information about these dedicated RCC channels can not be found on the status menu.

8.1.1 Channel IDs - Host Channel

Go to: View and Edit Channels > (Host Channel) > View and Edit SCSI ID

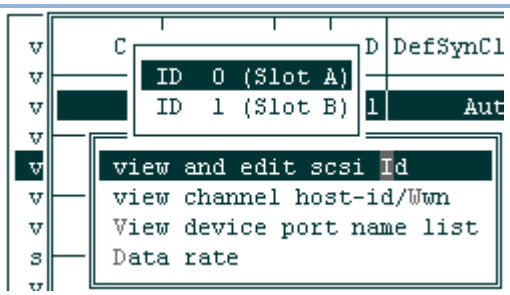

NOTE

If a host channel connection is configured in an arbitrated FC loop, in the Loop-only mode, the maximum number of host IDs per channel will be limited to “16.”

8.1.2 Adding an ID (Slot A / Slot B Controller ID)

Go to: View and Edit Channels > (Host Channel) > View and Edit SCSI ID > (ID) > Add Channel ID > (Slot) > (ID)

In a single-controller mode, the Slot B controller ID is unavailable. In a dual-controller configuration, you should manually create one or more Slot B controller IDs on your host channels.

The co-existing Slot A and Slot B IDs enable storage volumes to be presented through host ports on Slot A or Slot B controllers. You may refer to Chapter 16 of this manual for the configuration samples.

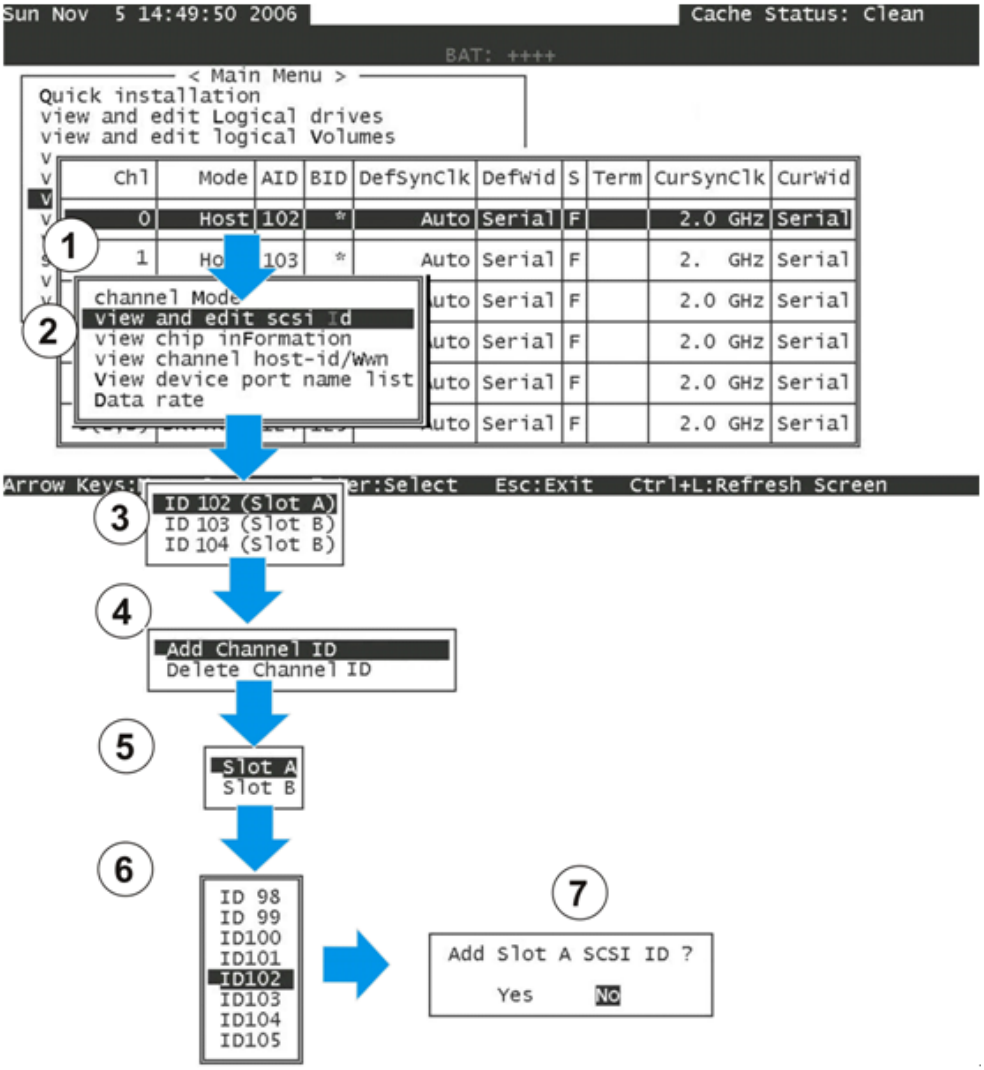

Once Slot B controller IDs are available, you can associate logical arrays with both Slot A and Slot B IDs so that system workload can be shared between partner controllers.

A confirm box will prompt reminding you that the configuration change will only take effect after the controller resets. Select Yes to confirm.

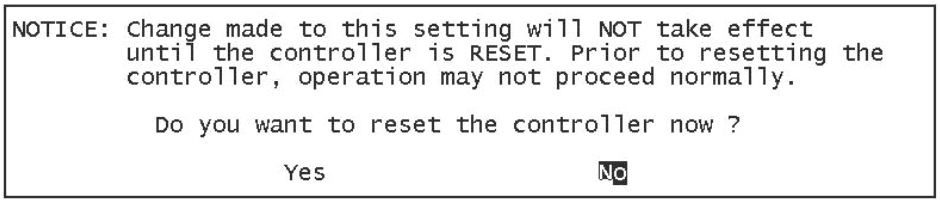
NOTICE: Change made to this setting will NOT take effect
until the controller is RESET. Prior to resetting the
controller, operation may not proceed normally.
Do you want to reset the controller now ?
Yes No

8.1.3 Deleting an ID

Go to: View and Edit Channels > (Host Channel) > View and Edit SCSI ID > (ID) > Delete Channel ID

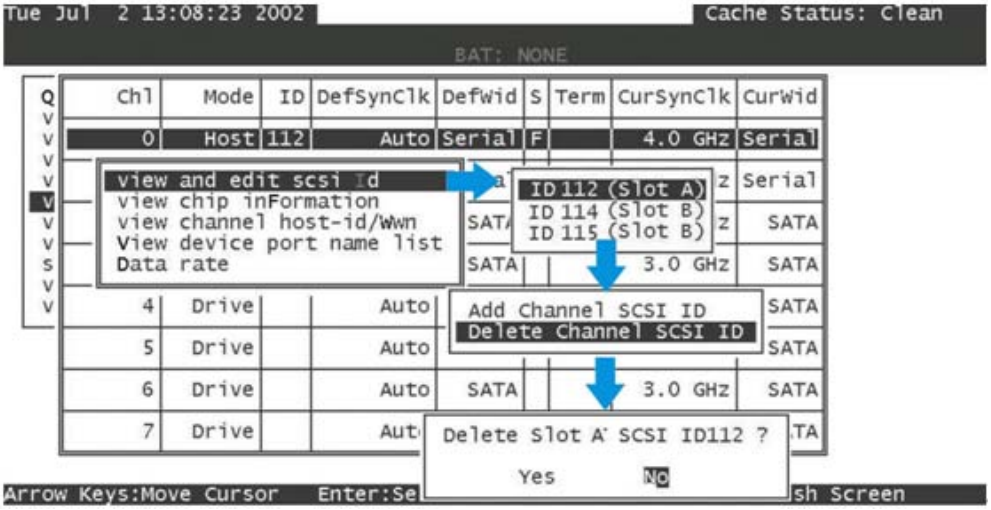
Tue Jul 2 13:08:23 2002
Cache Status: Clean
BAT: NONE
Chl Mode ID DefSyncClk DefWid S Term CurSyncClk CurWid
0 Host 112 Auto Serial F 4.0 GHz Serial
view and edit scsi Id
view chip inFormation
view channel host-id/Wwn
View device port name list
Data rate
ID 112 (Slot A)
ID 114 (Slot B)
ID 115 (Slot B)
Add Channel SCSI ID
Delete Channel SCSI ID
4 Drive Auto SATA
5 Drive Auto SATA
6 Drive Auto SATA 3.0 GHz SATA
7 Drive
Delete Slot A SCSI ID112 ?
Yes No
Arrow Keys:Move Cursor

A confirm box will prompt reminding you that the configuration change will only take effect after the controller resets. Select Yes to confirm.

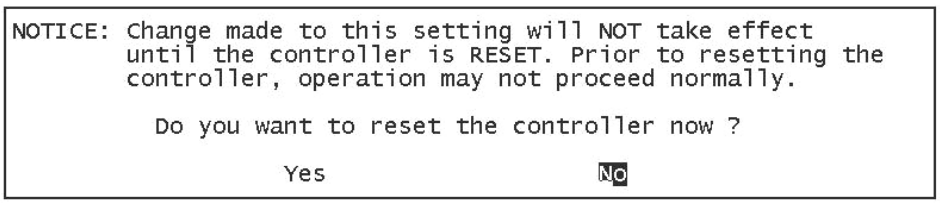
NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally.
Do you want to reset the controller now ?
Yes No

NOTE

Every time you add/delete a channel ID, you must reset the system for the changes to take effect.

Multiple target IDs can co-exist on a host channel while every drive channels in a dual-controller subsystem has two preset IDs.

At least one ID should be present on each channel bus.

For details on the relationship between host IDs and physical configurations in a dual-controller configuration, please refer to Chapter 16 Redundant Controller.

8.1.4 Setting Data Rate (Channel Bus)

Go to: View and Edit Channels > (Fibre Host/Drive Channel) > Data Rate

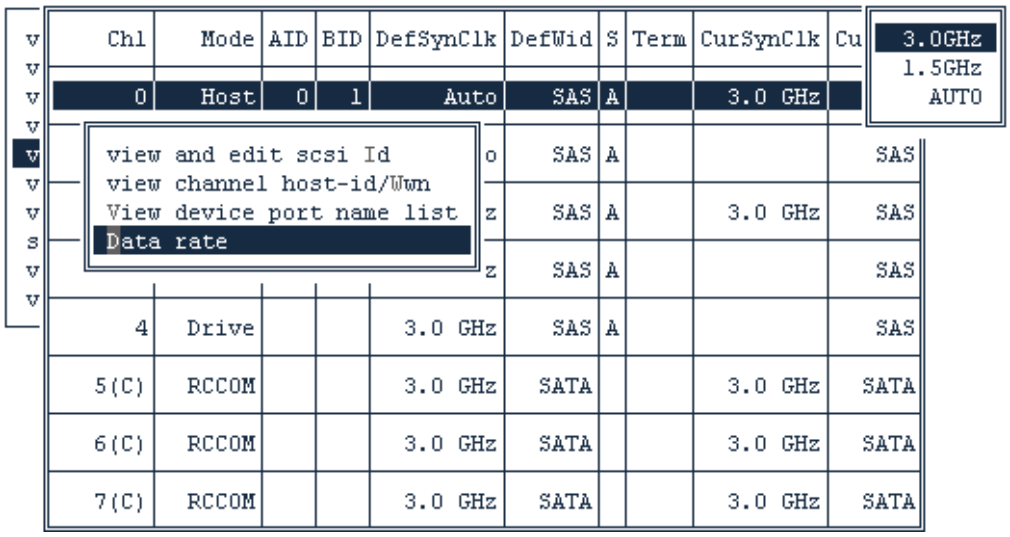
Chl Mode AID BID DefSynClk DefWid S Term CurSynClk Cu
0 Host 0 1 Auto SAS A 3.0 GHz
3.0GHz
1.5GHz
AUTO
view and edit scsi Id
view channel host-id/Wwn
View device port name list
Data rate
4 Drive 3.0 GHz SAS A SAS
5(C) RCCOM 3.0 GHz SATA 3.0 GHz SATA
6(C) RCCOM 3.0 GHz SATA 3.0 GHz SATA
7(C) RCCOM 3.0 GHz SATA 3.0 GHz SATA

This option is available in the configuration menu of Fibre host channel and of the drive channel configuration menus in Fibre-, SAS-, or SATA-based subsystems. Default is “AUTO” and should work fine with most disk drives. Changing this setting is not recommended unless some particular bus signal issues occur.

8.1.5 Viewing Channel Host ID/WWN

Go to: View and Edit Channels > (Host Channel) > View Channel Host-ID

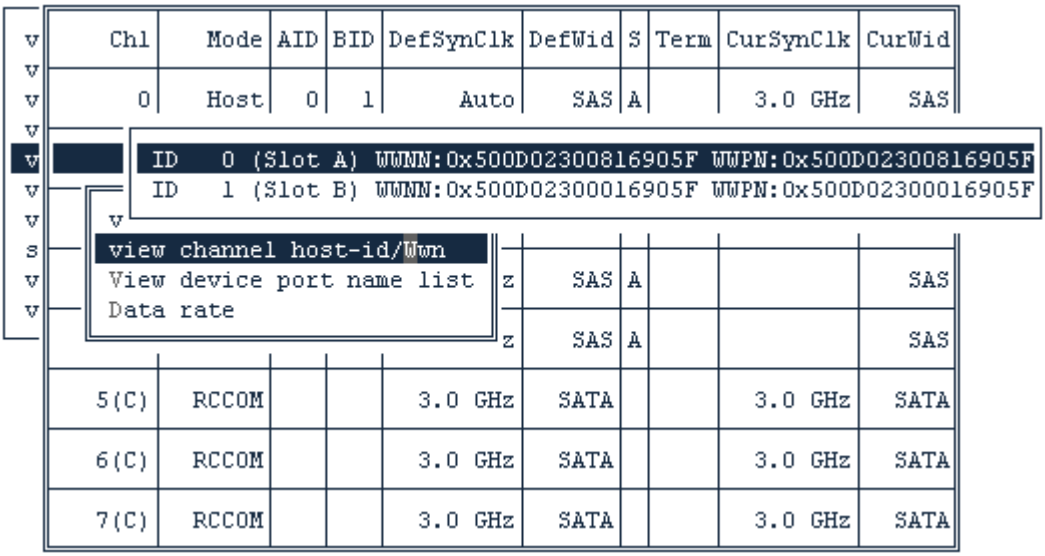
Chl Mode AID BID DefSynClk DefWid S Term CurSynClk CurWid
0 Host 0 1 Auto SAS A 3.0 GHz SAS
ID 0 (Slot A) WWNN:0x500D02300816905F WWPN:0x500D02300816905F
ID 1 (Slot B) WWNN:0x500D02300016905F WWPN:0x500D02300016905F
view channel host-id/Wwn
View device port name list
Data rate
5(C) RCCOM 3.0 GHz SATA 3.0 GHz SATA
6(C) RCCOM 3.0 GHz SATA 3.0 GHz SATA
7(C) RCCOM 3.0 GHz SATA 3.0 GHz SATA

WWNN Node name and WWPN Port name are unique eight-byte addresses that appear on a Fibre Channel host port. Every host channel ID appears as a Fibre Channel device and carries both a node name and a port name. A storage volume associated with a host ID will also be associated with both a unique node name and a port name.

Corresponding to the dual-ported connectivity defined in Fibre Channel specifications, some of the SAN management software on the market may identify a RAID storage by checking its specific node name and port names.

If a storage volume needs to appear through fault-tolerant links, it needs to be associated with host IDs on separate host ports (channels).

Two identical host IDs (e.g., ID0 on CH0 and ID0 on CH1) on two different host channels carry an identical node name. If an volume is associated with these IDs, the array will appear with one node name and two different port names. Some management software will then be able to identify these port names as alternate data paths to a storage device.

The Host-ID/WWN option allows users to inspect the node names and port names assigned to specific host IDs.

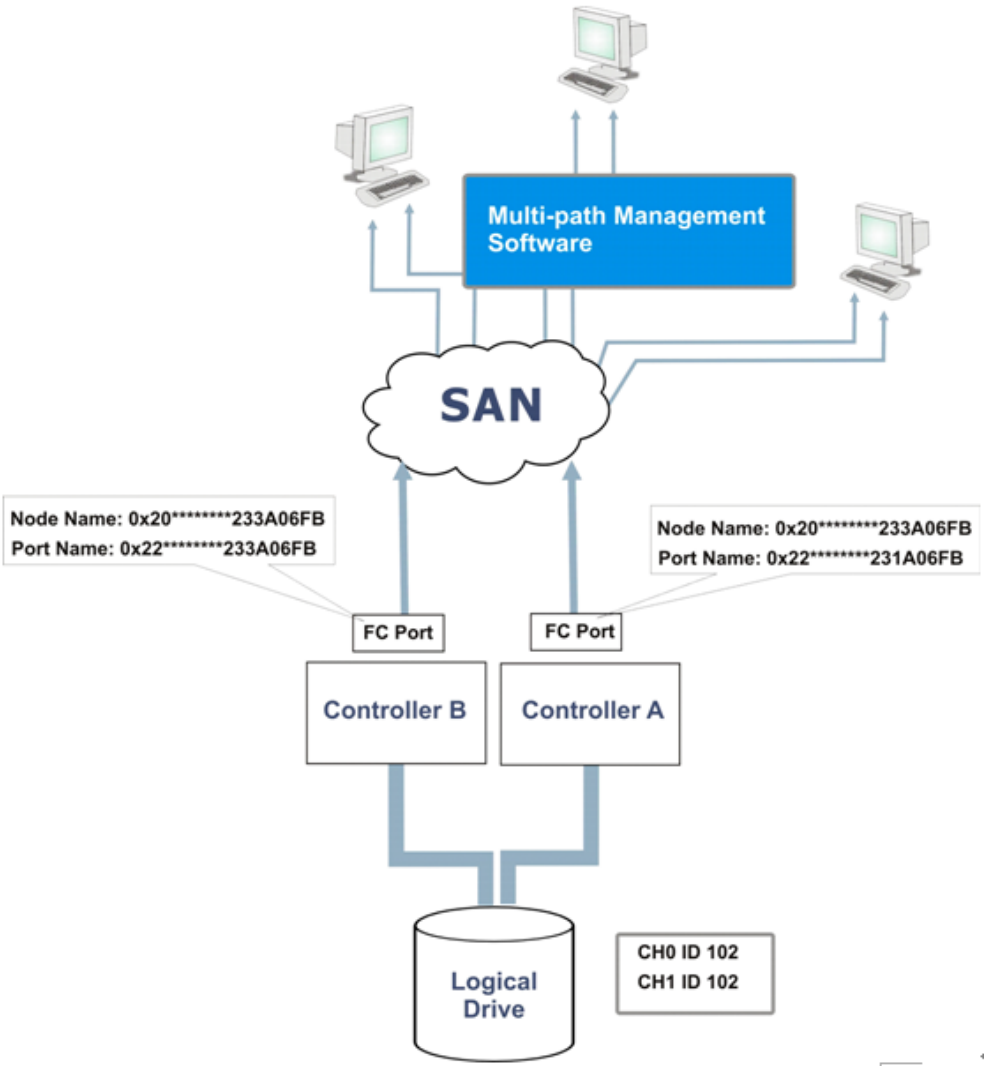

8.1.6 Viewing Device Port Name List (WWPN)

Go to: View and Edit Channels > (Device Channel)

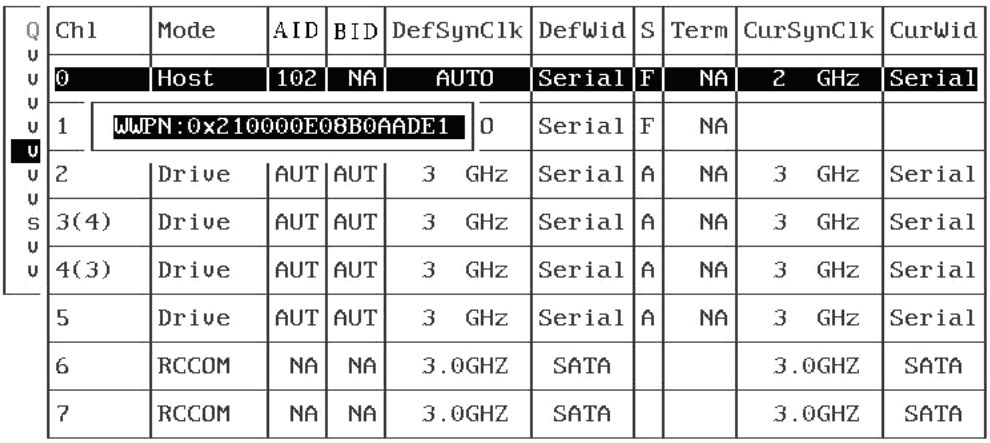
Chl Mode AID BID DefSynClk DefWid S Term CurSynClk CurWid
0 Host 102 NA AUTO Serial F NA 2 GHz Serial
1 WWPN:0x210000E08B0AADE1 0 Serial F NA
2 Drive AUT AUT 3 GHz Serial A NA 3 GHz Serial
3(4) Drive AUT AUT 3 GHz Serial A NA 3 GHz Serial
4(3) Drive AUT AUT 3 GHz Serial A NA 3 GHz Serial
5 Drive AUT AUT 3 GHz Serial A NA 3 GHz Serial
6 RCCOM NA NA 3.0GHZ SATA 3.0GHZ SATA
7 RCCOM NA NA 3.0GHZ SATA 3.0GHZ SATA

This function displays the device port names (host adapter WWN) of the adapters that appear on the connection with a host port or through a switched fabric connection.

The HBA port names detected here can be manually added to the "Host-ID WWN name list" in "View and Edit Host LUN" menu. Giving nicknames to HBA ports can ease identification of FC ports in SAN and facilitate LUN mapping processes.

8.1.7 Adding Host – ID/WWN Label Declaration

Go to: View and Edit Channels > (Device Channel)

A nickname can be appended to any host adapter WWN for ease of identification in SAN environments where multiple servers reside in a storage network. Choose Yes and enter a name for the host adapter port.

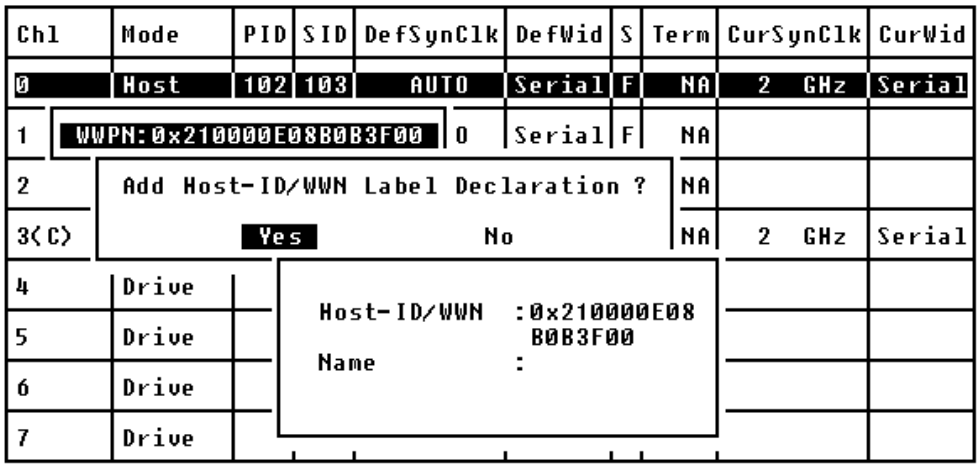
Chl Mode PID SID DefSynClk DefWid S Term CurSynClk CurWid
0 Host 102 103 AUTO Serial F NA 2 GHz Serial
1 WWPN:0x210000E08B0B3F00 0 Serial F NA
2 Add Host-ID/WWN Label Declaration ? NA
3(C) Yes No NA 2 GHz Serial
4 Drive
5 Drive
6 Drive
7 Drive
Host-ID/WWN :0x210000E08 B0B3F00
Name :

8.2 Setting Fibre-related Host-side Parameters

Go to: View and Edit Configuration Parameters > Host-side Parameters > Fibre Connection Option

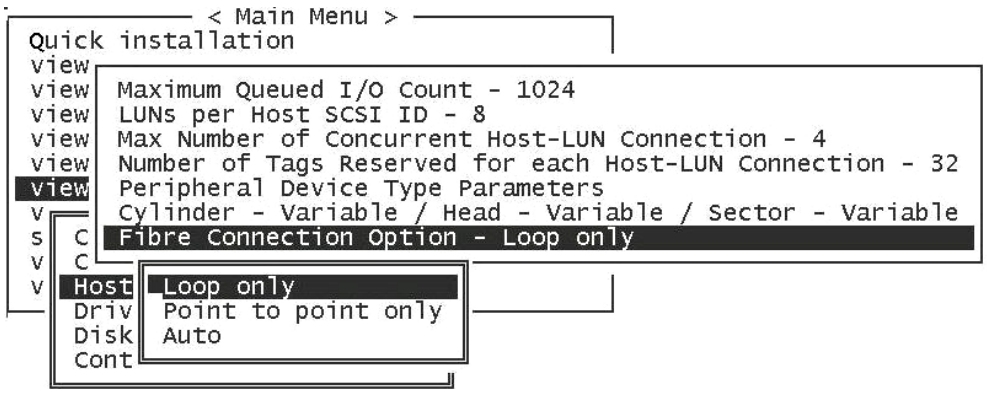

8.2.1 About Loop Only

Go to: View and Edit Configuration Parameters > Host-side Parameters > Fibre Connection Option > Loop Only

The firmware default is Loop only. Under the following conditions you may want to configure the Fibre connection type to Loop only:

You prefer using multiple host IDs on a single host channel for a complex SAN configuration. If set to Point-to-point, there is only one host ID for each Fibre channel on each controller.

Most FC switches can automatically detect port communication protocols. The Loop Only option will work for most configurations.

The below drawing presents a faulty configuration:

The connection points #1, #2, #3, and #4 form a public loop. However, a public loop does not allow two FL_ports on two different switches to co-exist. One of the FL_ports will fail.

8.2.2 About Point-to-point

Go to: View and Edit Configuration Parameters > Host-side Parameters > Fibre Connection Option > Point-to-Point

If you connect your Fibre-host Galaxy HDX4 system to SAN, you may consider

setting the host protocol to “Point-to-point.” However, doing so will limit the number of host channel IDs. You will then use LUN numbers under IDs for multiple mapping instances.

If host protocol is set to Point-to-Point, each controller can have a host ID (AID / BID) on each host channel, meaning that an AID will be available through the controller A FC port, and another BID through the controller B FC port. This enables high availability storage configuration.

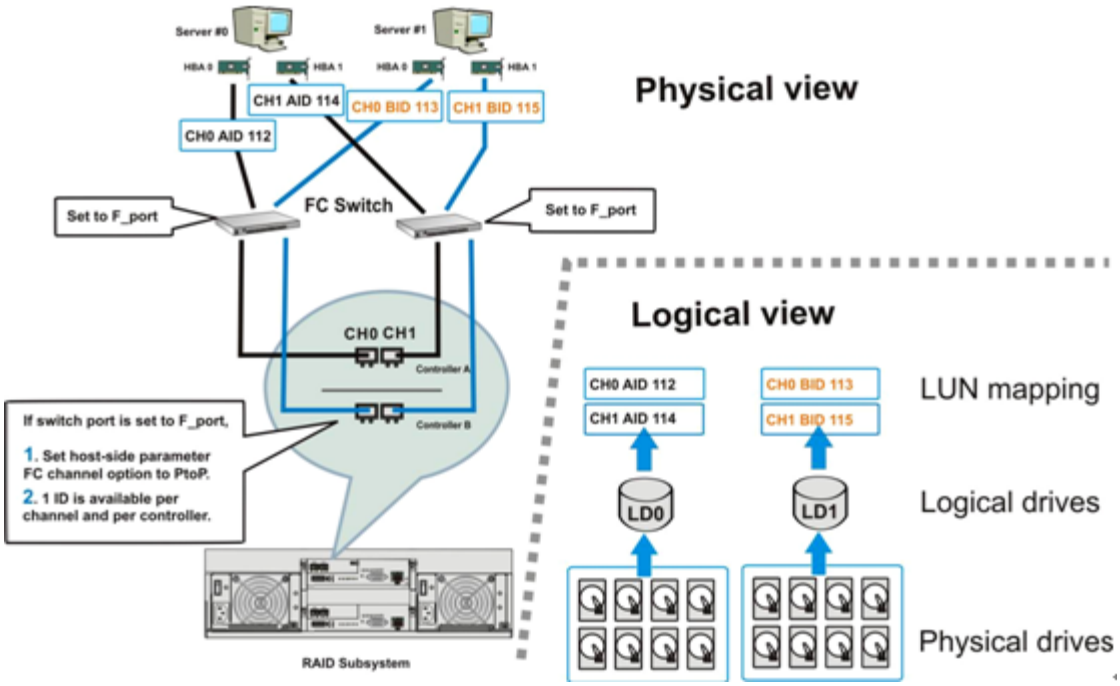

8.2.3 Setting Controller Unique Identifier

Go to: View and Edit Configuration Parameters > Controller Parameters > Controller Unique Identifier

NOTE

If you install a RAID controller to a system and for some reasons move it to another system, this controller may have already acquired a unique ID from the previous enclosure. Port names and MAC address conflicts will occur if you attach these two systems to a Fibre Channel storage network.

If the ID conflicts occur, you have to restore NVRAM defaults.

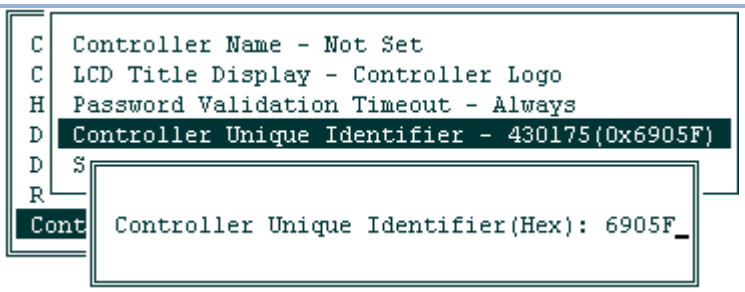
Controller Name - Not Set
LCD Title Display - Controller Logo
Password Validation Timeout - Always
Controller Unique Identifier - 430175(0x6905F)
Controller Unique Identifier(Hex): 6905F_

A Controller Unique Identifier is required for operation with the Redundant Controller Configuration. All Galaxy HDX4 subsystems come with a preset identifier.

The unique identifier will be used to generate a Fibre Channel "node name" (WWNN). The node name is device-unique and comprised of information such as the IEEE company ID and this user-configurable identifier in the last two bytes.

In redundant mode when a controller fails and a replacement is combined, the node name will be passed down to the replacement, making the host unaware of controller replacement so that the controller failover and failback process can complete in a host-transparent manner.

All Galaxy HDX4 subsystems come with a default identifier. This identifier guarantees your FC ports' port names and node names are unique over a Fibre Channel network. Making changes to the default value is only necessary if the port name conflicts should occur.

9 iSCSI Options

This chapter guides users through all configuration processes of an iSCSI storage system. Some features, e.g., Grouping (Multiple Connections per Session) and SLP (Service Location Protocol), require the mutual support from counterpart devices in an iSCSI network.

NOTE

If you apply Microsoft's software initiator and also Galaxy's multi-pathing driver, please do not select the multi-path checkbox while configuring target portals. There is a similar checkbox that will appear during the installation of the software initiators. It is recommended you deselect the checkbox.

The samples below start from simple to complex, load-sharing, redundant-controller configurations. Multiple logical drives will be created to be managed by partner RAID controllers depending on the number of RAID and JBOD enclosures.

9.1 iSCSI Topology Examples

9.1.1 iSCSI IP SAN – 1 Topology

Single controller model (G models)

# Host	# HBA	# Channel	# Cable	# Controller	# LD	# Map	Hub status
2	4 with 1	4	8	1	2	4	None

port

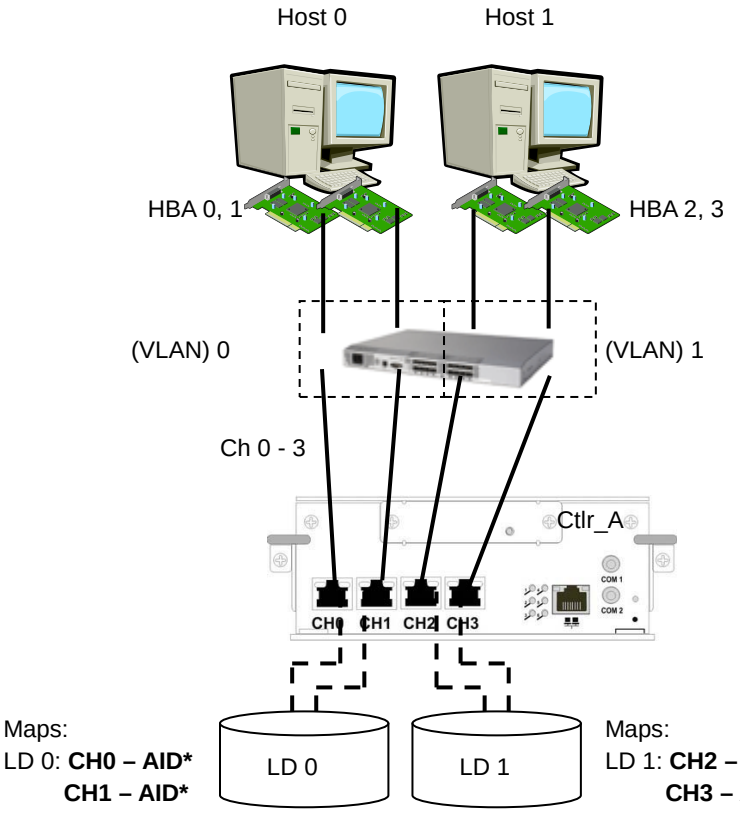

Type	Failed	Available Path	Multi-path
HBA failure	HBA 0	Host 0 -> LD 0: HBA 1 – VLAN0 – CH0 – AID – LD 0 Host 0 -> LD 0: HBA 1 – VLAN0 – CH1 – AID – LD 0	✓
One Cable Failure	HBA 0 – VLAN0	Same with host HBA 0 failed	✓
	VLAN0 - Ctlr_A	Host 0 -> LD 0: HBA 0 – VLAN0 – CH1 – AID – LD 0 Host 0 -> LD 0: HBA 1 – VLAN0 – CH1 – AID – LD 0	✓
Switch Failure	Zone 0	X (Host 0 -> LD 0)	
Controller failure	Ctlr_A	X (Host 0 -> LD 0)	
Controller absent	Ctlr_A	X	

9.1.2 iSCSI IP SAN – 2 Topology

Single controller model (G) for clustering environment

# Host	# HBA	# Channel	# Cable	# Controller	# LD	# Map	Hub status
2	4 with 1 port	4	8	1	2	4	None

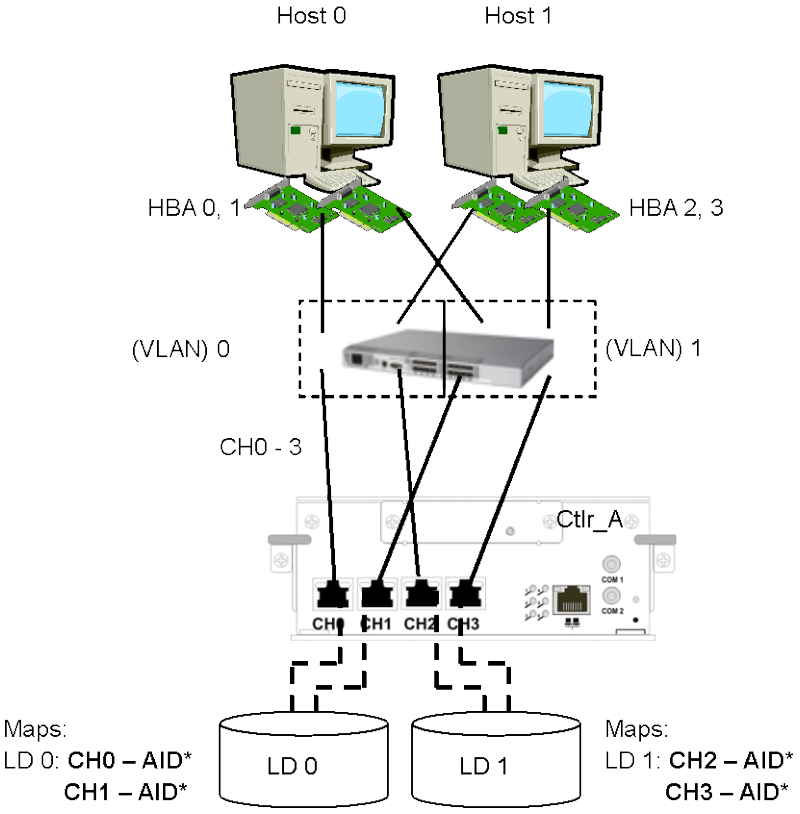

Type	Failed	Available Path	Multi-path
HBA failure	HBA 0	X (Failed over handled by clustering OS)	✓
One Cable Failure	HBA 0 – VLAN0	Same with host HBA 0 failed	✓
	VLAN0 - Ctlr_A	Host 0 -> LD 0: HBA 1 – VLAN1 – CH1 – AID – LD 0 Host 1 -> LD 0: HBA 3 – VLAN1 – CH1 – AID – LD 0	✓
Switch Failure	Zone 0	Same as cable in VLAN0 – Ctrl_A failed	✓
Controller failure	Ctlr_A	X	
Controller absent	Ctlr_A	X	

9.1.3 iSCSI IP SAN – 3 Topology

Redundant controller models, for one host, two hosts or clustering environment

# Host	# HBA	# Channel	# Cable	# Controller	# LD	# Map	Hub status
2 (1)	4 with 1 port	2	12 (10)	2	2	4	None

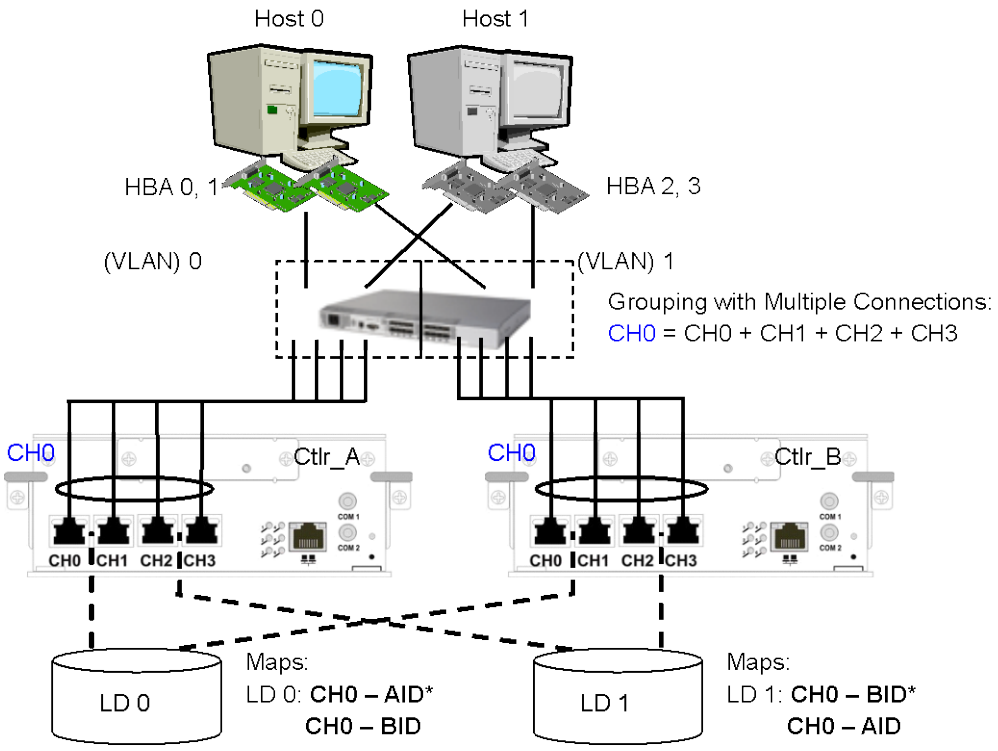

Type	Failed	Available Path	Multi-path
HBA failure	HBA 0	Host 0 -> LD 0: HBA 1 – VLAN1 – CH0 – BID – LD 0 (re-route) Host 0 -> LD 1: HBA 1 – VLAN1 – CH0 – BID – LD 1	✓
One Cable Failure	HBA 0 – VLAN0	Same with host HBA 0 failed	✓
	VLAN0 - CH0 - Ctlr_A	One cable failed: no effect, both cables failed: Host 0 -> LD 0: HBA 1 – VLAN1 – CH0 – BID – LD 0 (re-route) Host 0 -> LD 1: HBA 1 – VLAN1 – CH0 – BID – LD 1	✓
Switch Failure	VLAN0	Same as cable in VLAN0 – Ctrl_A failed	✓
Controller failure	Ctlr_A	Host 0 -> LD 0: HBA 1 – VLAN1 – CH0 – AID – LD 0 Host 0 -> LD 1: HBA 1 – VLAN1 – CH0 – BID – LD 1	✓
Controller absent	Ctlr_A	X	✓

NOTE: For a configuration with two application servers and without clustering, users should use LUN Masking or VLAN to avoid different hosts from accessing the same LD causing data contention.

9.1.4 iSCSI IP SAN – 4 Topology

Redundant controller models , for one host, two hosts or clustering environment

# Host	# HBA	# Channel	# Cable	# Controller	# LD	# Map	Hub status
2 (1)	4 with 1	4	12 (10)	2	4-8	8-16	None

port

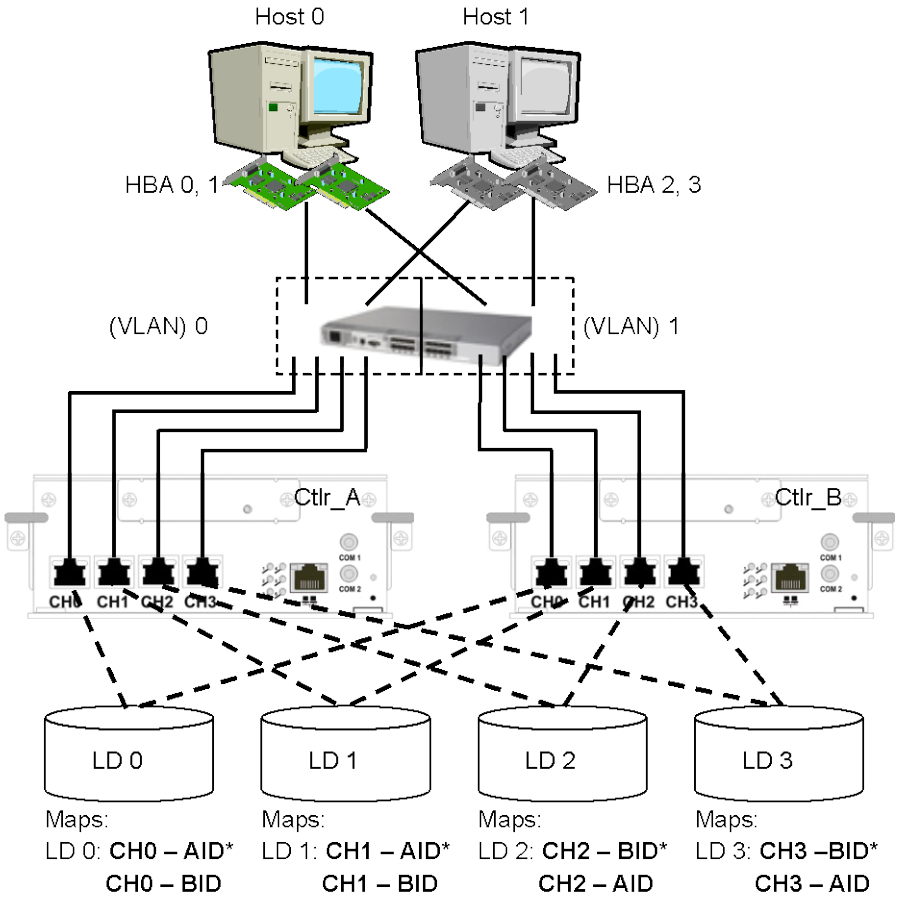

NOTE: For making the most out of this configuration, users can create 4 or more LDs (LD4, LD5, LD6, and LD7) and assign passive paths as their active data paths such as the following:

LD 4: **CH2 – AID*** **CH2 – BID** LD 5: **CH3 – AID*** **CH3 – BID** LD 6: **CH0 – BID*** **CH0 – AID** LD 7: **CH1 – BID*** **CH1 – AID**

Type	Failed	Available Path	Multi-path
HBA failure	HBA 0	Host 0 -> LD 0: HBA 1 – VLAN1 – CH0 – BID – LD 0 (re-route)	✓
		Host 0 -> LD 1: HBA 1 – VLAN1 – CH1 – BID – LD 1 (re-route)	
		Host 0 -> LD 2: HBA 1 – VLAN1 – CH2 – BID – LD 2	
		Host 0 -> LD 3: HBA 1 – VLAN1 – CH3 – BID – LD 3	
One Cable Failure	HBA 0 – VLAN0	Same with host HBA 0 failed	✓
	VLAN0 - CH0 - Ctlr_A	Host 0 -> LD0: HBA 1 – VLAN1 – CH0 – BID – LD 0 (re-route)	✓
		Host 0 -> LD1: HBA 0 – VLAN0 – CH1 – AID – LD 1	
		Host 0 -> LD2: HBA 1 – VLAN1 – CH2 – BID – LD 2	
		Host 0 -> LD3: HBA 1 – VLAN1 – CH3 – BID – LD 3	
Switch Failure	VLAN0	Same as cable in VLAN0 – Ctrl_A failed	✓
Controller failure	Ctlr_A	Host 0 -> LD0: HBA 1 – VLAN1 – CH0 – AID – LD 0	✓
		Host 0 -> LD1: HBA 1 – VLAN1 – CH1 – AID – LD 1	

		Host 0 -> LD2: HBA 1 – VLAN1 – CH2 – BID – LD 2 Host 0 -> LD3: HBA 1 – VLAN1 – CH3 – BID – LD 3	
Controller absent	Ctlr_A	Same as cable in VLAN0 – Ctrl_A failed	✓

NOTE: For a configuration consisting of multiple servers without clustering, users should use VLAN to avoid different host from accessing the same LD causing data contention.

9.2 Trunking

Trunking is implemented following IEEE standard 802.3. The “Trunk Group” function is available with this firmware revision.

Use Limitations:

- Correspondence with Channel MC/S group (see the following section on grouping):
  Because of the order in protocol layer implementation,
  You cannot configure MC/S grouped channels into trunks.
  Yet you can configure trunked ports into MC/S groups.
- Channel IDs: If multiple host ports are trunked, IDs will be available as if on one channel.
- IP Address Setting: Trunked ports will have one IP address. Trunked ports reside in the same subnet.
- LUN Mapping: LUN mapping to a trunked group of ports is performed as if mapping to a single host port.
- Switch Setting: The corresponding trunk setting on switch ports should also be configured, and it is recommended to configure switch setting before changing system setting.

9.2.1 Setting Switch Trunk Port

Sample pages of switch trunk port settings (3COM 2924-SFP Plus) are shown below:

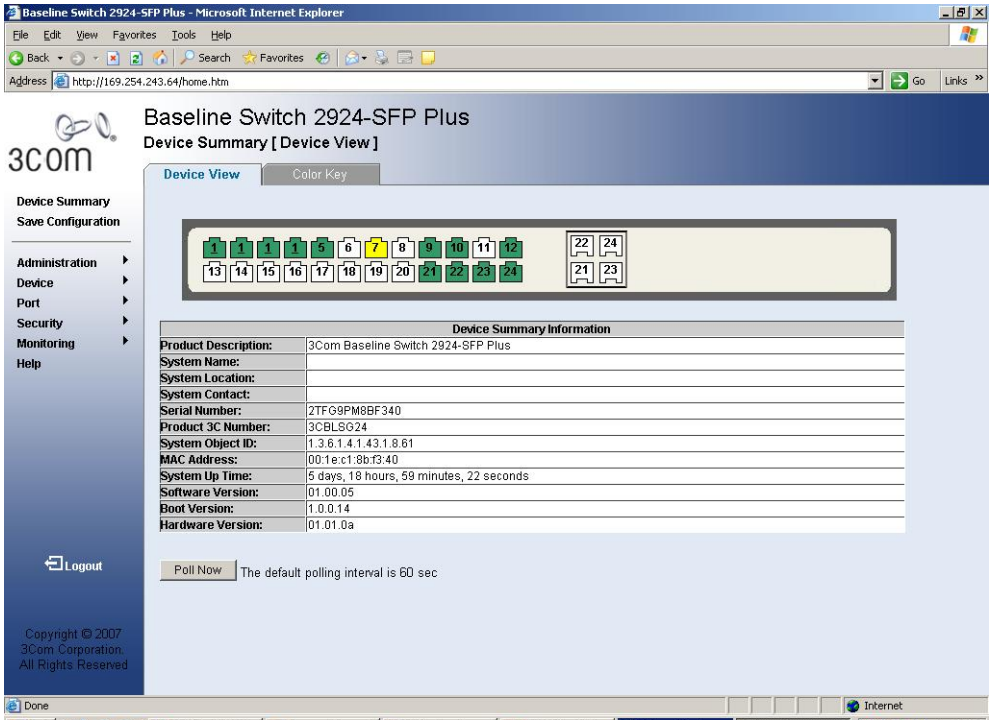

Configuration is done via Port -> Link Aggregation -> Aggregation group ID.

Port selection is done via LACP -> Select port.

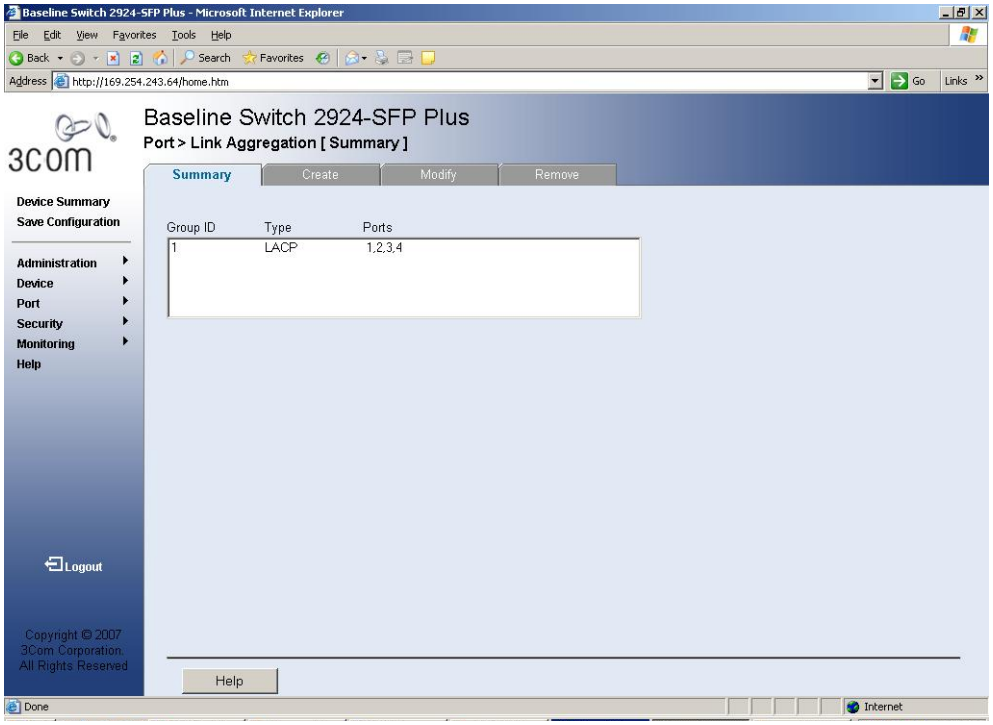

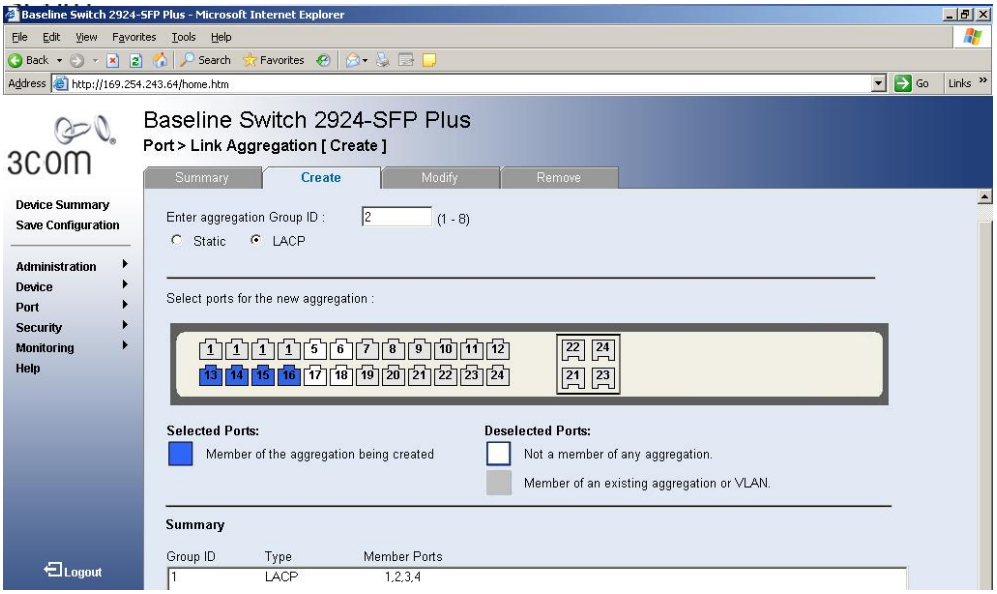

Refer to the documentation that came with your Ethernet switches for instructions on trunk port configuration.

Make sure you have appropriate configurations both on your iSCSI system and Ethernet switches. Otherwise, networking failures will occur.

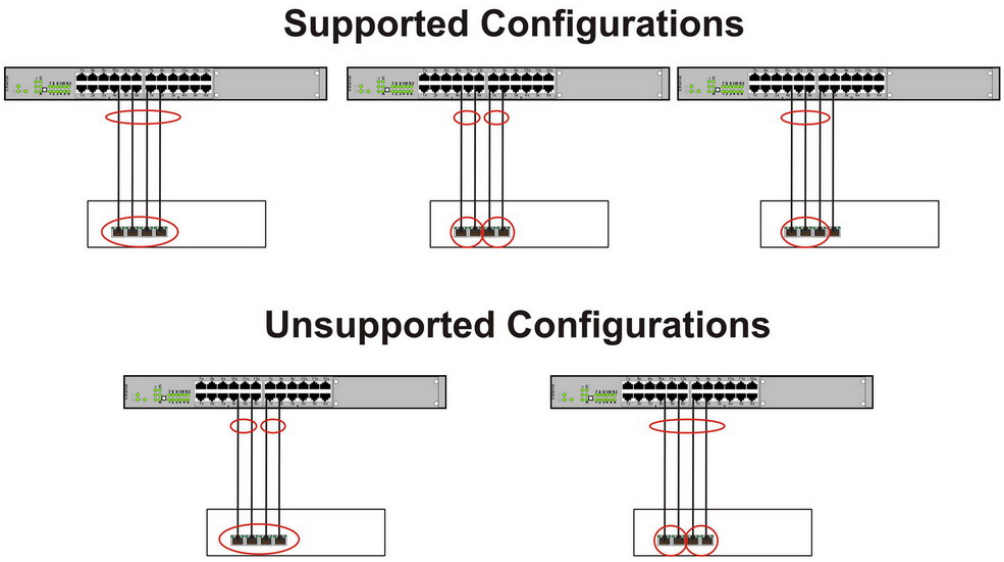

9.2.2 Notes on Trunking Conditions

- Aggregation interfaces must be connected in the same network, often the same Ethernet switch, limiting the physical isolation of the multiple paths.
- Trunking implementation is dependent on having aggregation-capable devices and switches.

- All ports can be trunked into a single IP, or several IPs. For example, there are 4 GbE ports in iSCSI storage system and user can configure those 4 ports into a single IP, or two IPs each by trunking two physical ports. Trunked ports combinations can be 2, 3, 4, 5, 6.
- After trunking is complete, you need to reset the storage subsystem.
- If a trunk configuration is not valid, firmware will report a trunk failure event. For example, with 4 GbE ports into a trunk on an iSCSI storage system, while the corresponding ports on GbE switch are not trunked. If so, the trunking configuration is not completed and another event will prompt. Users should configure switch settings and reboot iSCSI storage system again.
- Requirements on system reset after making changes to trunk configuration:
  Create new trunk groups or change member ports
  Change trunk group ID
  Change IP address: Reset (as usual, both iSCSI host ports and the 10/100BaseT mgmt. port)
- Trunking and iSCSI MC/S (Multiple Connections per Session):
  Configure port trunking before MC/S configuration.
  If there are any configured MC/S groups when creating IP trunking, remove those MC/S groups.

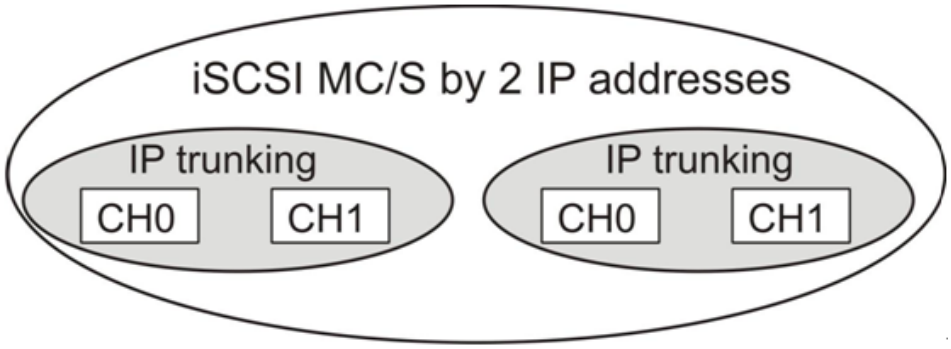

- Link Aggregation, according to IEEE 802.3, does not support the following:
  Multipoint Aggregations: The mechanisms specified in this clause do not support aggregations among more than two systems.
  Dissimilar MACs: Link Aggregation is supported only on links using the IEEE 802.3 MAC (Gigabit Ethernet and FDDI are not supported in parallel but dissimilar PHYs such as copper and fiber are supported)
  Half duplex operation: Link Aggregation is supported only on point-to-point links with MACs operating in full duplex mode.
  Operation across multiple data rates: All links in a Link Aggregation Group

operate at the same data rate (e.g. 10 Mb/s, 100 Mb/s, or 1000 Mb/s).

- Users cannot remove a master trunk port from a trunk configuration, for example, CH0 of a trunk group consisting of channels 0, 1, 2, and 3. The first port (having a smallest index number) within a trunk group is considered a master port member. To break master port from the trunk group, you can delete the whole trunk group.

9.2.3 Configuring Trunk

Go to: View and Edit Configuration Parameters > Communication Parameters > View and Edit Trunk Group Setting

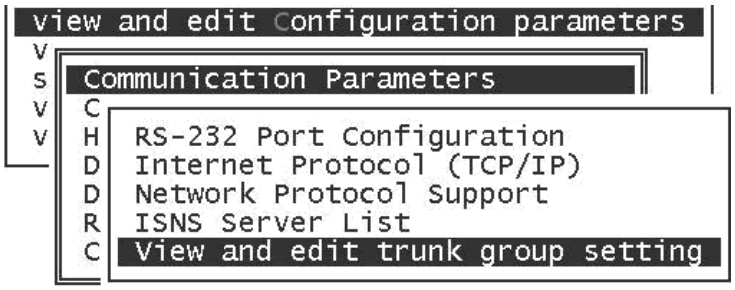

1. When prompted by Create Trunk Group?, press Enter on Yes to proceed.

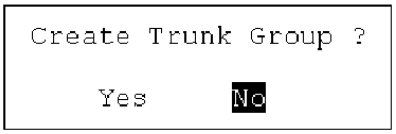

If there is no available host ports for trunk setting, or MC/S groups have been created, you will receive an error message saying “No available channel!”. If you have pre-configured MC/S groups, remove them before creating trunks.

2. Press Enter once to select each channel. Move to next channel using the arrow keys. Press ESC when you finish your selection.

CHL	IPv4	IPv6
1 0	10.10.10.10	Disable
2 1	10.10.10.20	Disable
2	10.10.10.30	Disable
3	10.10.10.40	Disable

3. There are channels that CANNOT be selected:
   Channels that have LUN mapping on them.
   Channels that are already trunked.
   Channels that are already included in MC/S groups.

4. Select Yes and press Enter on the confirm box.

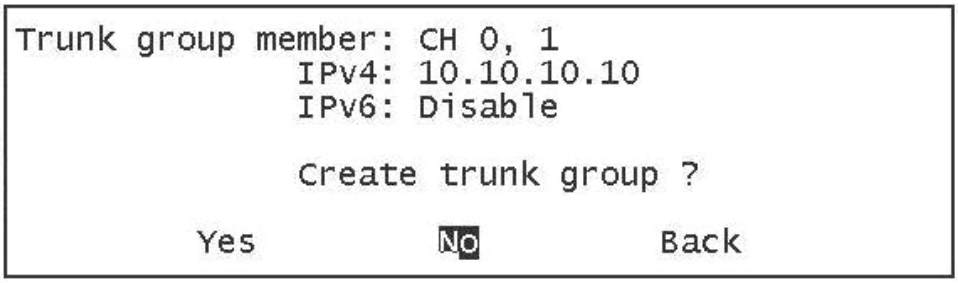

5. Trunked ports configuration may look like this. You can select up to 4 host ports into a trunk.

Member Chl	IPv4	IPv6
0, 1	10.10.10.10	Disable
2, 3	10.10.10.30	Disable

6. You can remove a member from a trunk group, or delete an existing group using the following commands.

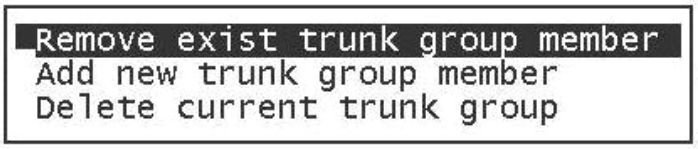

7. Note that you cannot remove a member if you have LUN mapping on the trunked ports.

8. Reset your iSCSI system for trunk setting to take effect.
   If your switch ports have not been configured, you will receive an error message saying trunk port configuration failure.
   If you configure ports 0 and 1 into trunk 1, and ports 2 and 3 into trunk 2, in View and Edit Channels menu, you can see that the corresponding channels are automatically configured into MC/S groups.

9.3 Grouping (MC/S, Multiple Connections per Session)

You can skip this section if you have already configured port trunking on your iSCSI host ports. Corresponding MC/S groups will be automatically created along with trunks.

9.3.1 Grouping VS Trunking

Grouping is different from Trunking. Trunking binds multiple physical interfaces so they are treated as one, and is accomplished in the TCP/IP stack. MC/S on the other hand allows the initiator portals and target portals to communicate in a coordinated manner. MC/S provides sophisticated error handling such that a failed link is recovered quickly by other good connections in the same session. MC/S is part of the iSCSI protocol that is implemented underneath SCSI and on top of TCP/IP.

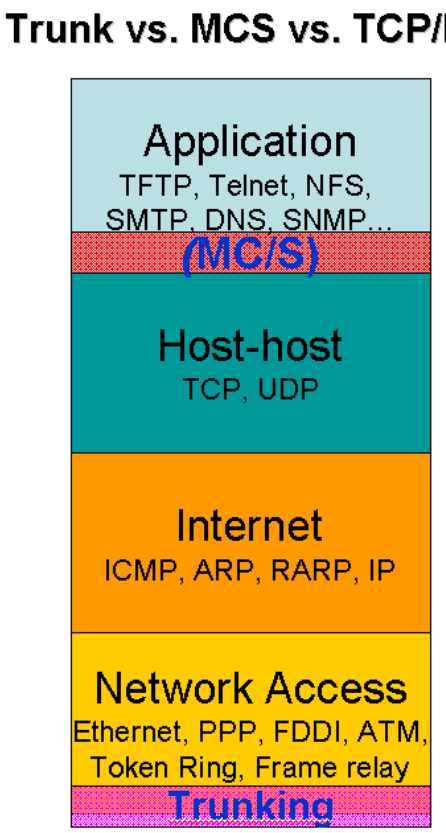

Grouping (MC/S) combines multiple host ports into a logical initiator-target session. MC/S can improve the throughput and transfer efficiency over a TCP session. Besides, this feature saves you the effort of mapping a logical drive to multiple host channel IDs on multiple host ports. The below drawings show 4 channels configured into an MC/S group.

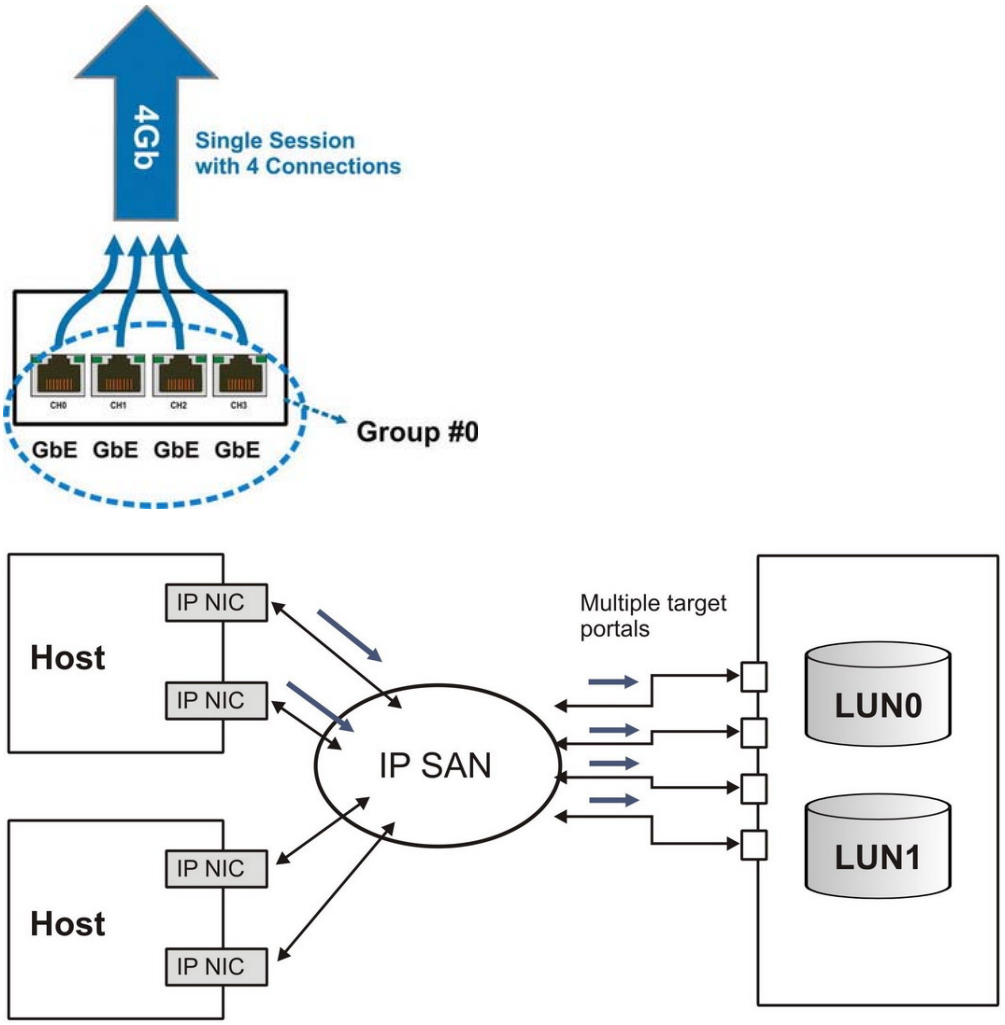

NOTE

If you prefer grouping (MC/S) using iSCSI TOE HBA cards, your HBAs must also support MC/S. Host initiators determine how I/O traffic is distributed through multiple target portals. Theoretically, I/Os are evenly distributed across physical paths.

9.3.2 Configuring Group

Go to: View and Edit Channels > (Host Channel) > Add to Group

Change the Group number on multiple channels to put them into the same logical group.

Chl	Mode	Grp	ID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	0	0	Auto	iSCSI			100MHz	iSCSI
1	Host	0	0	Auto	iSCSI			100MHz	iSCSI
2	Host	1	0	Auto	iSCSI			100MHz	iSCSI
3	Host	3	0	Auto	iSCSI			100MHz	iSCSI
				Auto	SATA				SATA
				Auto	SATA				SATA
6	Drive			Auto	SATA				SATA
7	Drive			Auto	SATA			1.5 GHz	SATA

view and edit scsi Id
view chip inFormation
add to Group

> NOTE
>
> Changing group configuration requires resetting your system.

With logical grouping, a logical drive mapped to a channel group will appear as one device on multiple data paths. This is very similar to the use of multi-pathing drivers.

9.3.3 LUN Presentation with and without Grouping

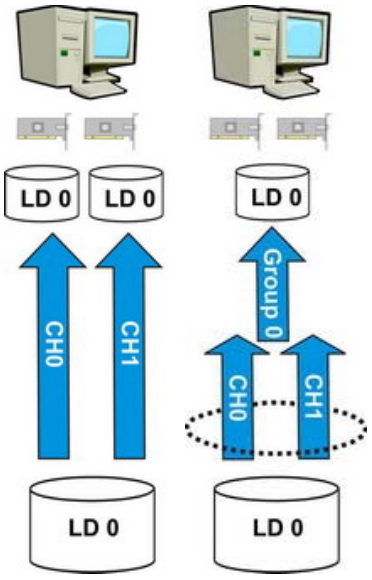

(Left figure) Grouping allows a consistent look of a storage volume to be seen over multiple connections, in a way very similar to the use of multi-pathing software. (Right figure) Without grouping, a storage volume will appear as two devices on two data paths.

> NOTE

For a redundant-controller system, you still need the multi-pathing driver to manage the data paths from different RAID controllers.

Appropriate configuration on software initiator is also necessary with grouping. The configuration process will be discussed later.

Host ports on different RAID controllers (a redundant-controller system) ARE NOT grouped together. Namely, in the event of a single controller failure, the IPs do not failover to the surviving controller.

9.3.4 Channels Automatically Divided into A and B Sub-groups

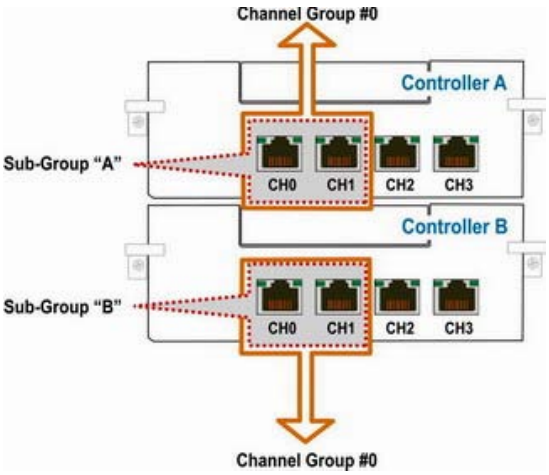

Chl	Mode	Grp	AID	BID	DefSyncClk	DefWid	S	Term	CurSyncClk	CurWid
0	Host	0	0	1	Auto	iSCSI			1GHz	iSCSI
1	Host	0			Auto	iSCSI			1GHz	iSCSI
2	Host	0			Auto	iSCSI			1GHz	iSCSI
3	Host	0			Auto	iSCSI			1GHz	iSCSI
4	Drive				3.0 GHz	SAS	A		3.0 GHz	SAS
5	Drive				3.0 GHz	SAS	A			SAS
6	Drive				3.0 GHz	SAS	A			SAS
7(C)	RCCOM				3.0 GHz	SATA			3.0 GHz	SATA

A parallel configuration logic is applied in Galaxy's firmware utility. On the firmware screen of a redundant-controller system, you can only see the channel settings for a single controller, yet actually they are automatically applied to the partner controller. If you configure "channel groups," you actually create juxtaposed groups on partner controllers. (see drawing above)

One volume mapped to both an AID and a BID will appear as two devices on the A

links and on the B links. You will then need the multi-pathing driver to manage the fault-tolerant paths.

9.3.5 LUN Presentation on Multiple Data Paths

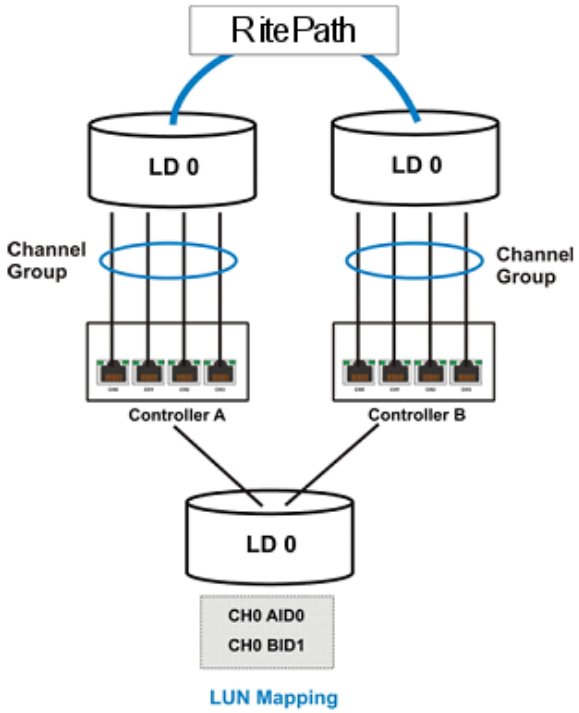

Here is a sample of 1 logical drive appearing as 2 devices across 8 data links (on 2 channel groups). With the help of RitePath, mapping to both controllers' IDs can ensure continuous access to data in the event of cabling or controller failure.

NOTE

Once channels are grouped, the channel group will behave as one logical channel, and the attributes of individual host channels will disappear. For example, if 4 channels are grouped together, only the IDs on the first channel remain.

Before Grouping		After Grouping	
Channel 0	ID 0	Channel 0	ID 0
Channel 1	ID 0	Channel 1	ID 0
Channel 2	ID 0	Channel 2	ID 0
Channel 3	ID 0	Channel 3	ID 0

Chl	Mode	Grp	AID	BID	DefSynClk	DefWid
0	Host	0	0	1	Auto	iSCSI
1	Host	0	-	-	Auto	iSCSI
2	Host	0	-	-	Auto	iSCSI
3	Host	0	-	-	Auto	iSCSI
4	Drive				3.0 GHz	SAS
5	Drive				3.0 GHz	SAS
6	Drive				3.0 GHz	SAS
7(C)	RCCOM				3.0 GHz	SATA

NOTE

Although the individual channel information is not available, you still need to take care of the TCP/IP connections. For example, you will need to consult your network administrator and configure a static port IP for your iSCSI host ports. The individual host port information is found under “View and Edit Configuration Parameters” -> “Communication Parameters” -> “Internet Protocol (TCP/IP)” -> “chx[LAN] MACAddr xxxxx.”

9.4 Setting Network Interface

9.4.1 IP Addresses to the iSCSI Host Ports

Go to: View and Edit Configuration Parameters > Communication Parameters > Internet Protocol

Contact your network administrator to obtain a list of valid IP addresses. Provide the adequate NetMask and Gateway values accordingly so that the host ports can connect to the initiators on your application servers.

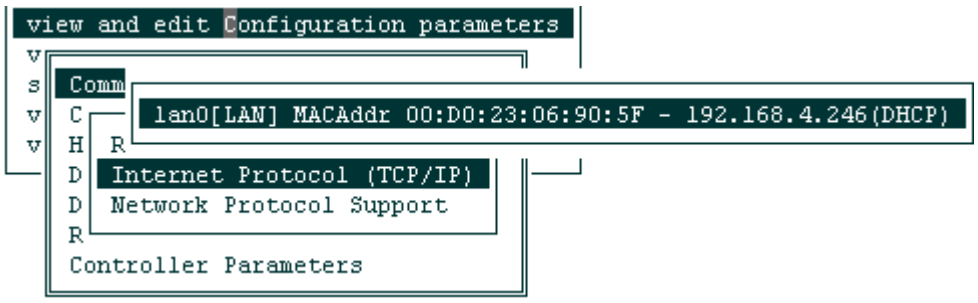

On a redundant-controller system each host channel is interfaced through 2 host

ports. One on controller A, and the other on controller B.

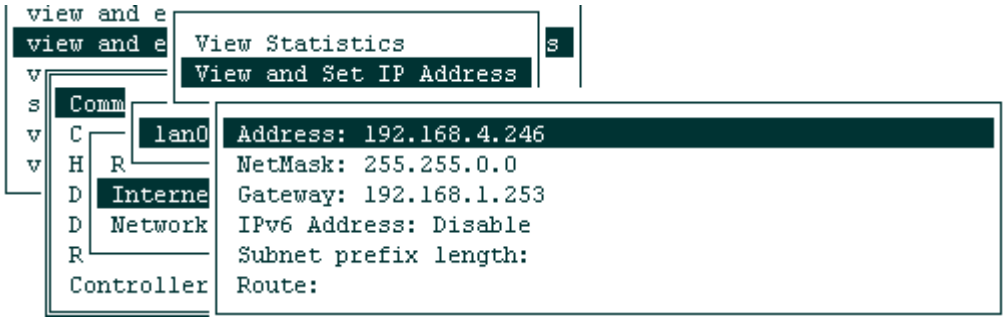

1. Press [ENTER] on a host port you wish to configure. The identity of a host port is presented as: “Channel number [LAN] MAC address – IP address (IP acquisition method)” The Slot A and Slot B addresses refer to the host ports on different RAID controllers.

> NOTE
>
> “lan0” is a 10/100BaseT management port for telnet or Galaxy Array Manager access.
>
> The corresponding controller A ports and controller B ports should be in the same subnet. In the event of controller failure, host should be able to access the alternate data paths on a surviving controller.

2. Press [ENTER] to select “Set IP Address”.

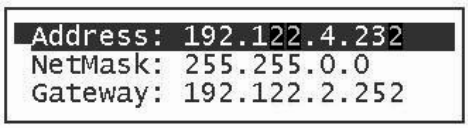

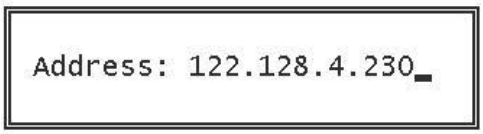

3. Reset the controller later when you finish all network settings.

9.4.2 Creating Host Channel IDs

Go to: View and Edit Channels > View and Edit SCSI ID > (ID) > Add Channel > (ID)

The redundant-controller system comes defaulted with 2 channel IDs, AID0 and

BID1, which are not sufficient for a more complex redundant controller settings. Create more IDs under the View and Edit Channel menu. Adding IDs requires rebooting your storage system.

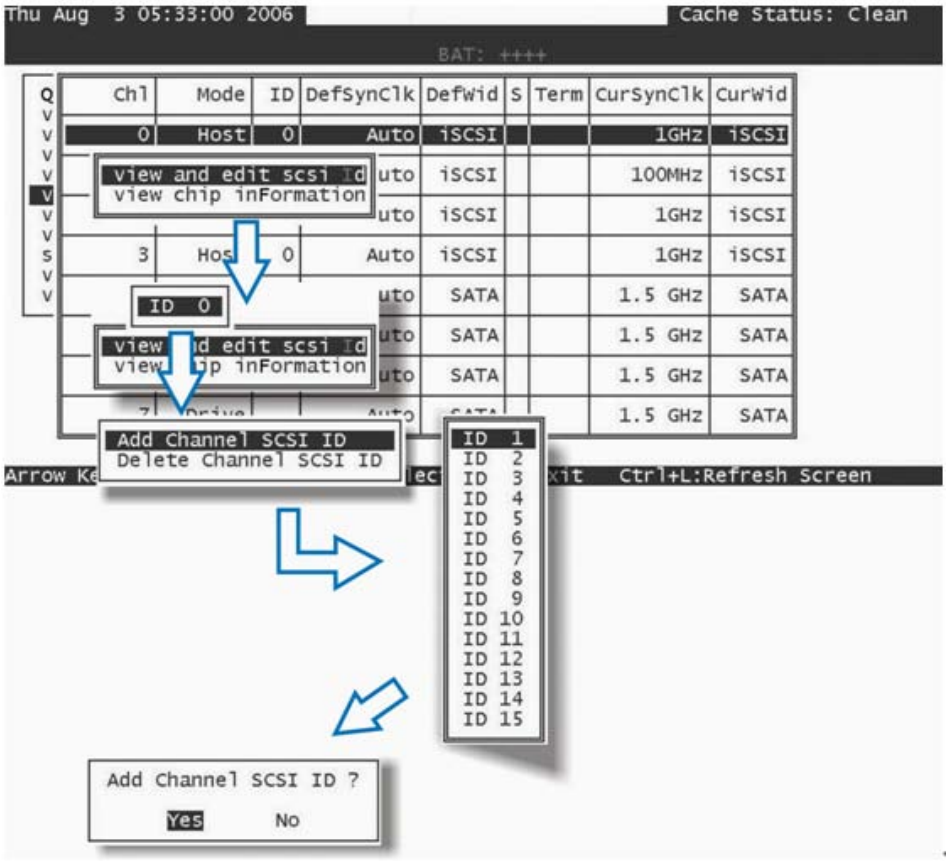

9.5 Mapping Host LUN

Once your network and RAID volume settings are done, install and enable initiators on your application servers. You can now turn on the network devices, storage system, and servers, and map your storage volumes to host LUNs so that network connectivity can be verified.

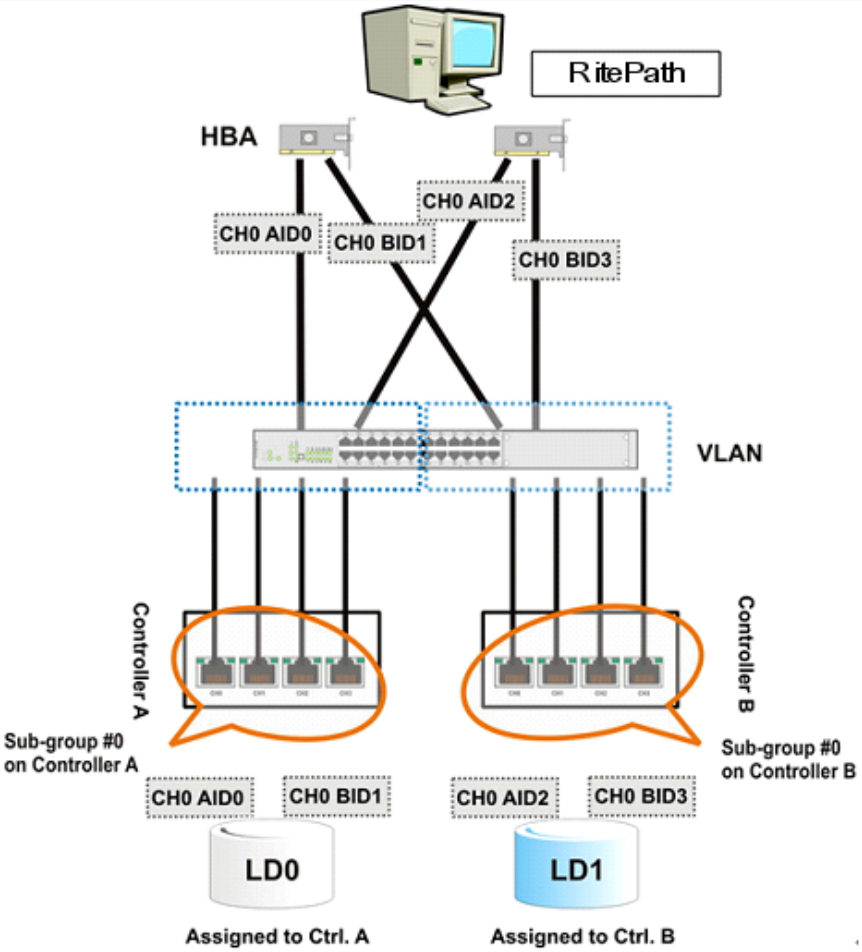

The above drawing shows a basic fault-tolerant configuration where service can continue with any single point of cabling or RAID controller failure. For simplicity reasons, only 1 server and 4 host links from it are shown. More logical drives, HBAs, or servers can attach to the configuration.

9.5.1 Creating an iSCSI Initiator List

Go to: View and Edit Host LUNs > Edit iSCSI Initiator List > Host

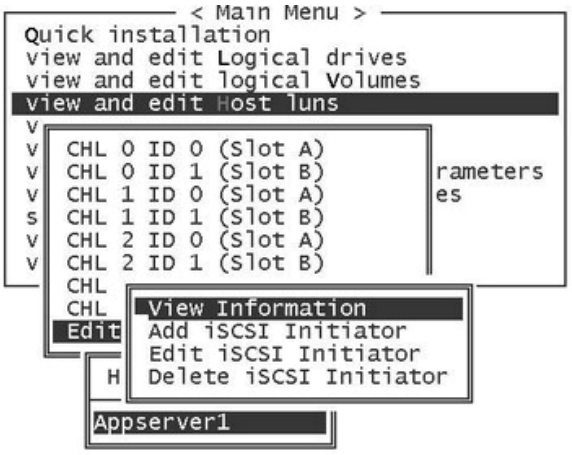

When your application servers are powered on, you should be able to see initiators from the firmware screen. Use the initiator list to organize your iSCSI connections.

In the initiator's attribute window, you can configure the following:

1. Enter a nickname for an initiator.
2. Manually key in an initiator's IQN.
3. Select an initiator from a list of IQNs that have already been detected.
4. Configure either the One-way or Mutual CHAP authentication (see 9.8. CHAP Login Authentication).
5. Apply IP Address and NetMask settings (if necessary). Multiple initiator ports on an application server can sometimes share the same IQN.

< Main Menu >
Quick installation
Host Alias Name - Appserver2
Host IQN - iqn.1991-05.com.microsoft:dt9305
User Name - Not Set
User Password - Not Set
Target Name - Not Set
Target Password - Not Set
IP Address - Not Set
NetMask - 255.255.255.0
Appserver2
Edit iSCSI Initiator List

Having a list of initiators in firmware can facilitate the process for configuring host LUN mapping and LUN Masking control. The initiator list also contains input entries for CHAP settings. See the following section for details.

9.5.2 Configuring Initiator (Using Microsoft Software Initiator)

In a redundant-controller iSCSI storage configuration, special attention should be paid to the configuration with host-side initiators. The following procedure will be exemplified using Microsoft iSCSI initiators.

First, you should jot down a list of host port IPs from the "View and Edit Channels" menu. In this sample procedure, there are 8 host port IPs.

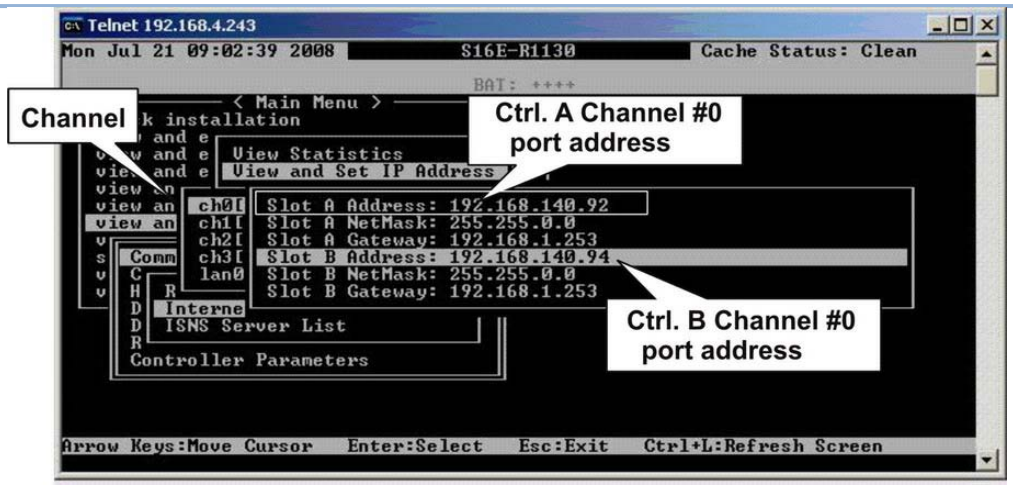

You may then develop a connection view diagram as follows:

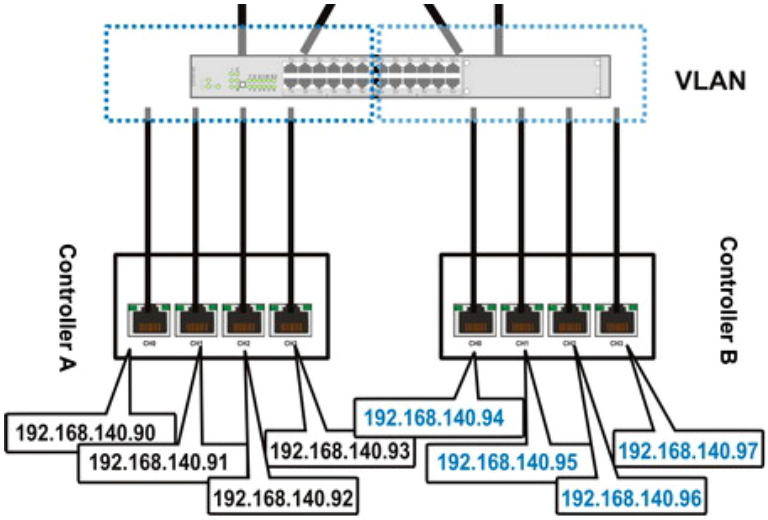

9.5.3 About IQN Name

There are several places you can see the IQN (iSCSI Qualified Name) of a logical drive (target) appearing through the network. IQN names are necessary when configuring iSCSI session or if you are using iSCSI TOE HBAs.

- You can scan the storage devices and see its IQN name through the initiator HBA utility or initiator software running on the host side, e.g., Microsoft iSCSI initiator.
- You can use the LCD keypad to find the serial number of a system. Find it in “Main Menu” –> “System Information” -> “Serial Number”

Galaxy’s storage IQN is composed of the system serial number and another 3 digits. The IQN always looks like the following:

iqn.2002-10.com.Galaxy:raid.snXXXXXX.XXX

- The 6 digits following the “sn” is the system’s serial number.
- The last 3 digits show variables in the following order:: “channel number” - “host ID” - “LD ownership”

The LD ownership digit shows either “1” or “2:” where “1” indicates Controller A and “2” indicates the LD ownership by the Controller B. Controller A is by default the dominating Primary controller.

The IQN is in accordance with how you map your logical drive to the host ID/LUN. For example, if you map a logical drive to host channel 0 and AID1, the last 3 digits will be 011.

9.5.4 Sample IQN Procedure

1. To configure multiple-portal access with the Grouping methodology, first open the initiator interface and manually key in a target port address, e.g., 192.168.140.90. Click the Add button to add a target port IP.

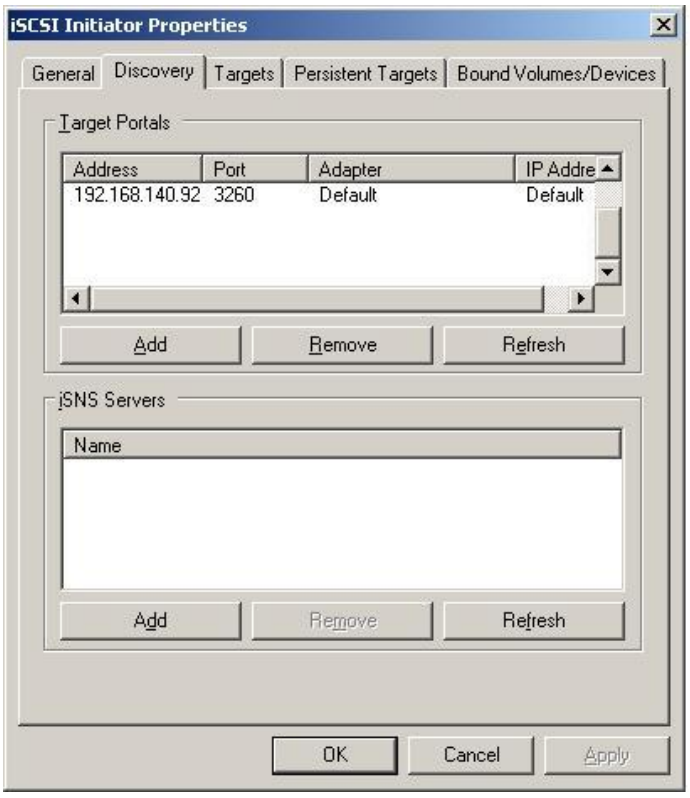

2. Under the Targets tab, select an IQN number from the list. From here you can identify the iSCSI targets by the last 3 digits of their IQN names.
   If the last digit is “1,” the target is a logical drive managed by controller A.
   If the last digit is “2,” the target is a logical drive managed by controller B.

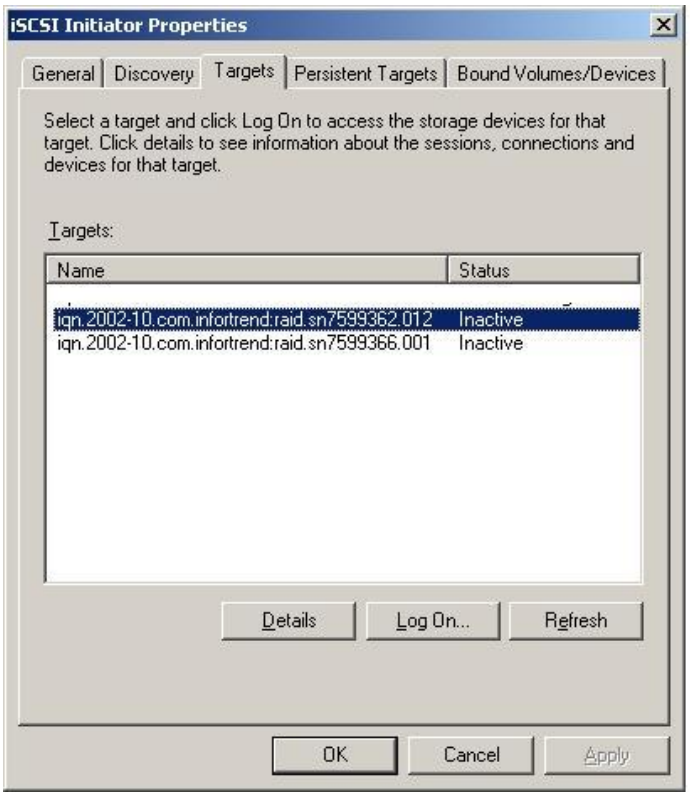

3. Click Log On… The Log On to Target window will prompt. Click the Advanced button on it.

4. Select the first check box, Automatically restore this connection when the system boots.

5. Select appropriate options for Local adapter, Source IP, and Target Portal from their respective pull-down lists.
   When selecting a Target Portal from the pull-down list, make sure you correctly associate a Target with the target portals. For example, a Target (Logical Drive) managed by Controller A should be associated with target portals that are controller A ports.
   Iqn.2002-10.com.Galaxy:raid.snXXXXXX.XX1 <- with -> A port target portals

Iqn.2002-10.com.Galaxy:raid.snXXXXXX.XX2 <- with -> B port target portals

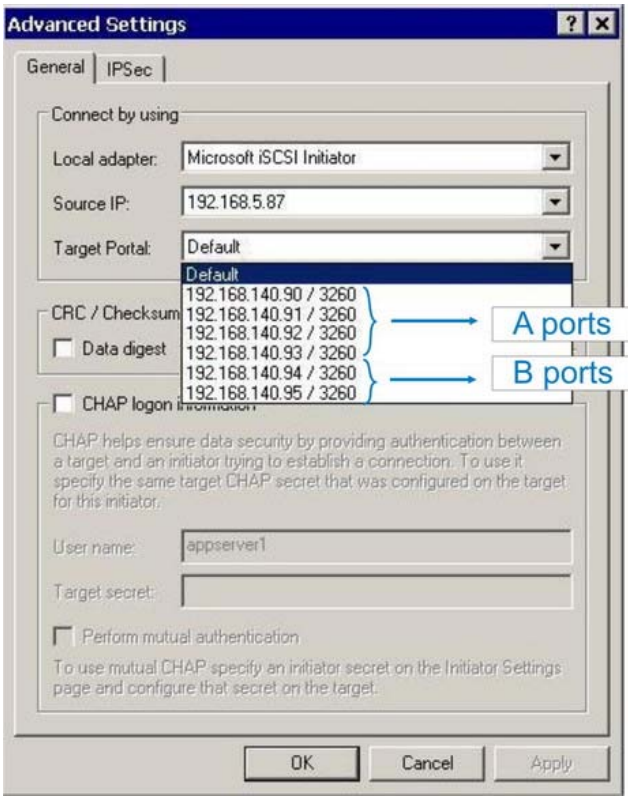

Click OK to close the window.

6. Return to the Targets window. Click on the Details button to add more target ports to this iSCSI target. The Target Properties window will prompt. Click the Connections button.

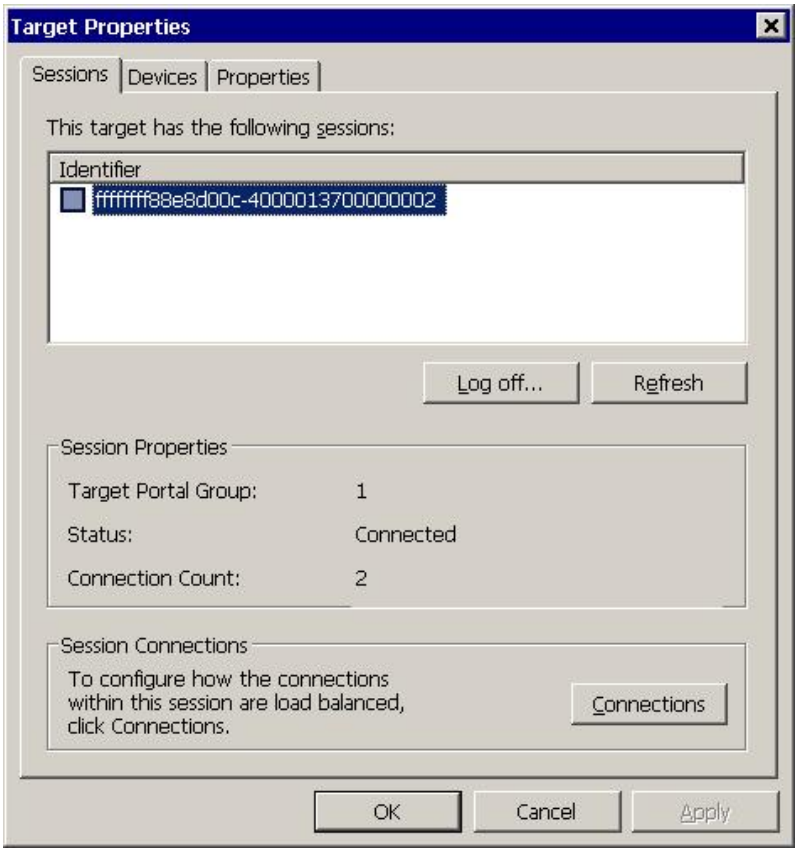

7. On the Session Connections window, select the Least Queue Depth load-balancing policy from the pull-down list, and then use the Add button below to include other target portals (A port portals at this stage) into the iSCSI session.

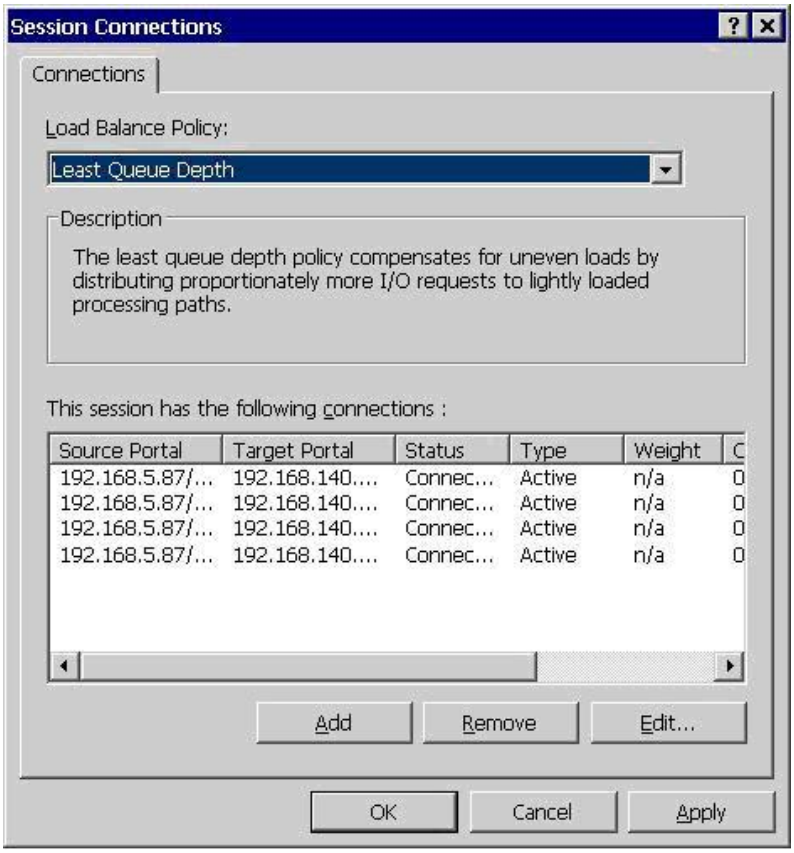

8. Click the Advanced button on the Add Connection box.

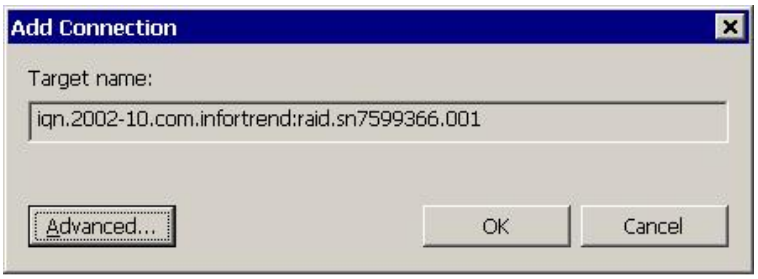

9. The Session Connections window will prompt again. Select a load-balancing policy and add a target port using the Add button. Click OK on the following screens to complete the configuration process.

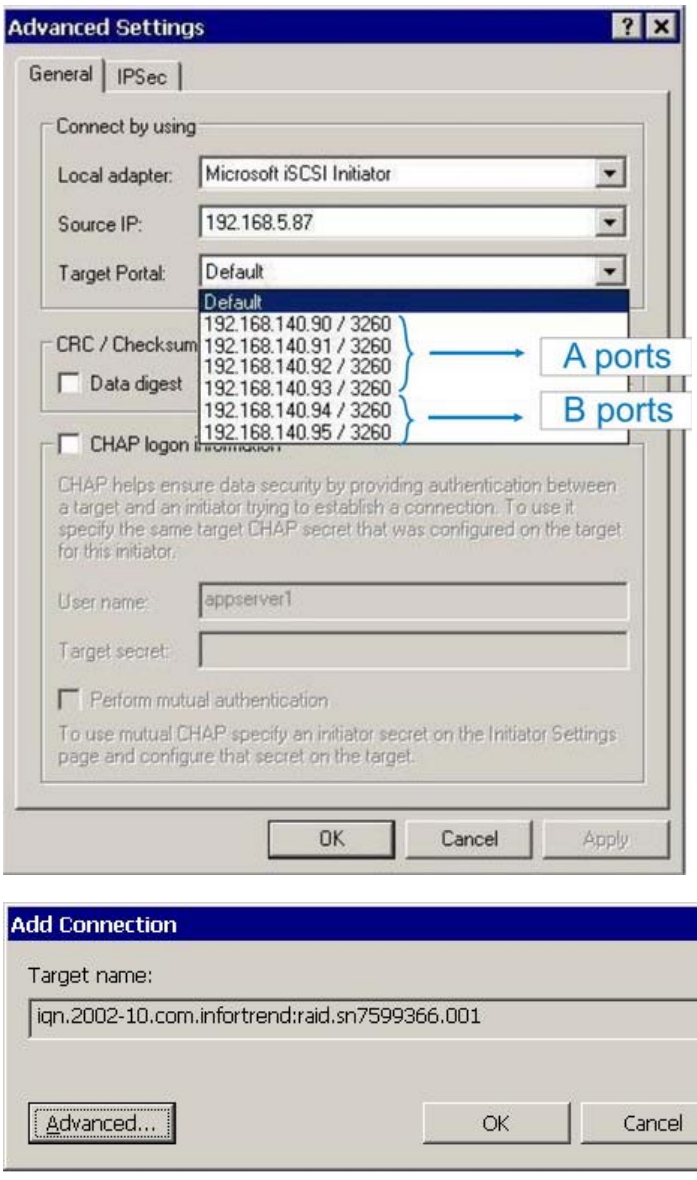

10. Repeat steps 6 through 9 until you add all portals available to the specific target.

11. Now you have finished configuring a logical drive target with target portals from controller A. You may then repeat the process above to associate another logical drive target with target portals from controller B. For example, you can start from adding 192.168.140.94 (a controller B port) in the Discovery window. Other portals from controller B, 192.168.140.95, …96, …97, will automatically appear in the pull-down list.

12. When you finish the configuration process, the logical drives will appear as multiple disk devices in the Windows Disk Drive management window. Install the RitePath multi-pathing software so that host can recognize them as devices accessed through fault-tolerant links.

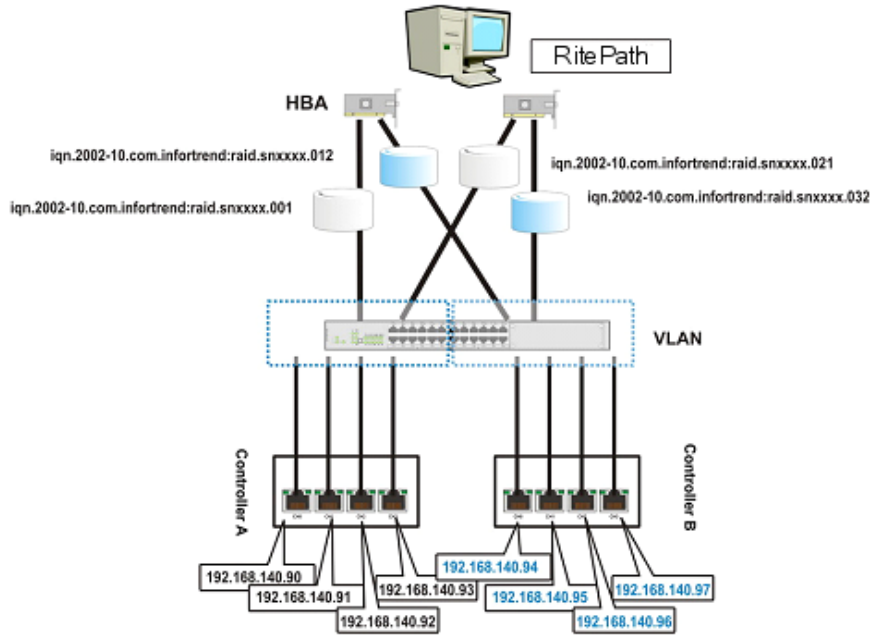

NOTE

Your iSCSI configuration may involve multiple servers and many logical drives. The maximum number of TCP sessions is 64 if you have a 1GB data cache in RAID controllers; and 32 sessions if using 512MB cache.

9.6 CHAP Login Authentication

NOTE

It is presumed that initiator software has already been installed and running on your application servers.

CHAP is one of the ways to authenticate access from networked servers to the iSCSI storage.

CHAP stands for Challenge Handshake Authentication protocol. With this protocol, the host-side initiators and storage systems use the encrypted password to authenticate each other remotely.

9.6.1 Configuring CHAP on RAID

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Login Authentication with CHAP

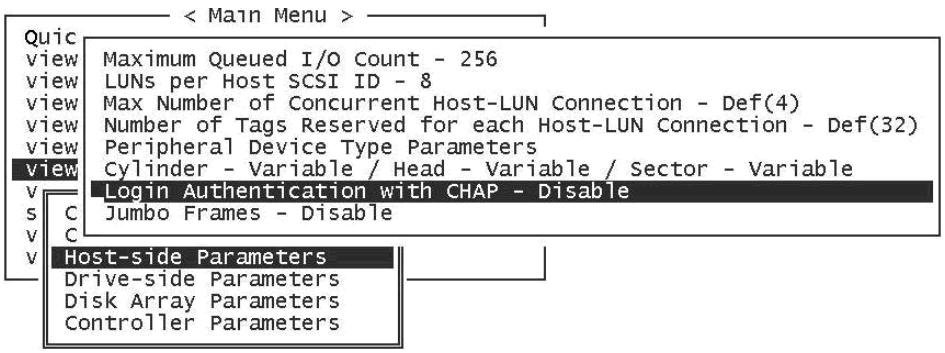
< Main Menu >
Maximum Queued I/O Count - 256
LUNs per Host SCSI ID - 8
Max Number of Concurrent Host-LUN Connection - Def(4)
Number of Tags Reserved for each Host-LUN Connection - Def(32)
Peripheral Device Type Parameters
Cylinder - Variable / Head - Variable / Sector - Variable
Login Authentication with CHAP - Disable
Jumbo Frames - Disable
Host-side Parameters
Drive-side Parameters
Disk Array Parameters
Controller Parameters

Both One-way and Two-way (mutual) CHAP authentications are supported. With Two-way CHAP, a separate three-way handshake is initiated between an iSCSI initiator and the storage system's host ports.

The CHAP-related options are found with the iSCSI Initiator List in your firmware screen under the "View and Edit Host LUNs" menu.

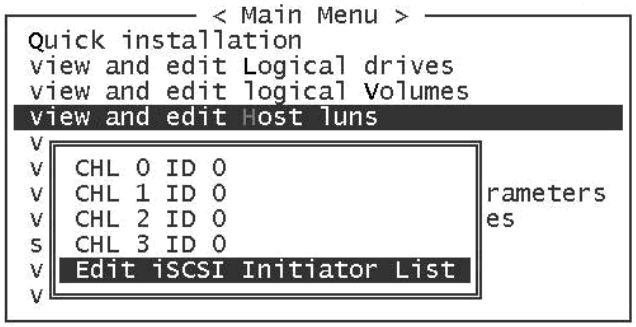
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
CHL 0 ID 0
CHL 1 ID 0
CHL 2 ID 0
CHL 3 ID 0
Edit iSCSI Initiator List

CHAP should be set by every initiator. Select and edit an initiator's attributes as shown in the following screen. The User Name, User Password, Target Name, and Target Password are used for CHAP authentication.

The User Password (One-way, from initiator) has to be at least 12 bytes, and the Target Password (Two-way, outbound from storage) has to be at least 14 bytes.

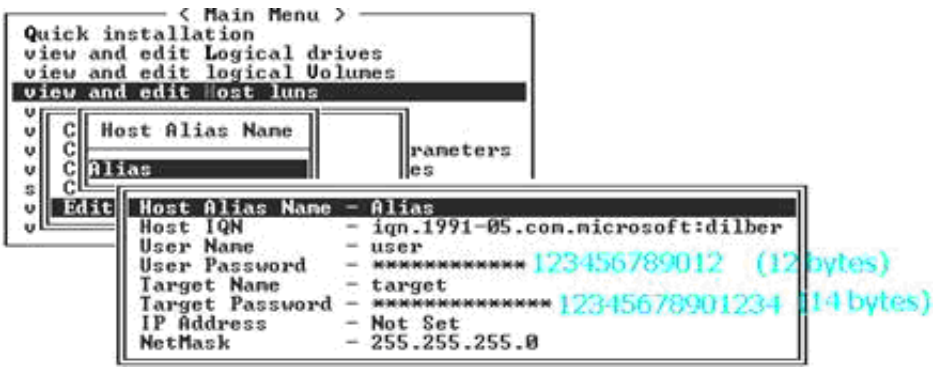
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
Host Alias Name
Alias
Host Alias Name - Alias
Host IQN - iqn.1991-05.com.microsoft:dilber
User Name - user
User Password - ********** 123456789012 (12 bytes)
Target Name - target
Target Password - ********** 12345678901234 (14 bytes)
IP Address - Not Set
NetMask - 255.255.255.0

The correspondence between firmware configuration entries and that of MS initiator is shown below:

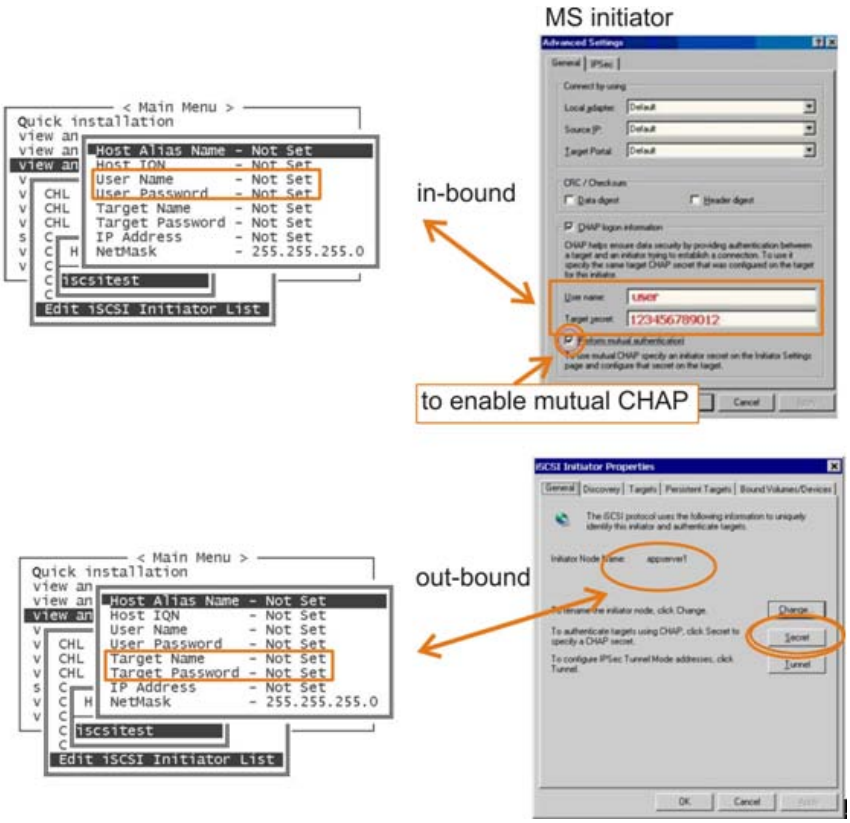

Enter identical names and passwords (secret) both in the firmware configuration screen above and in the initiators' screens below.

9.6.2 Configuring CHAP on the Initiator

1. Return to the Microsoft software initiator screen, under the General tab, enter the password for Two-way CHAP by clicking the Secret button (One that you entered in firmware for the Target Password).

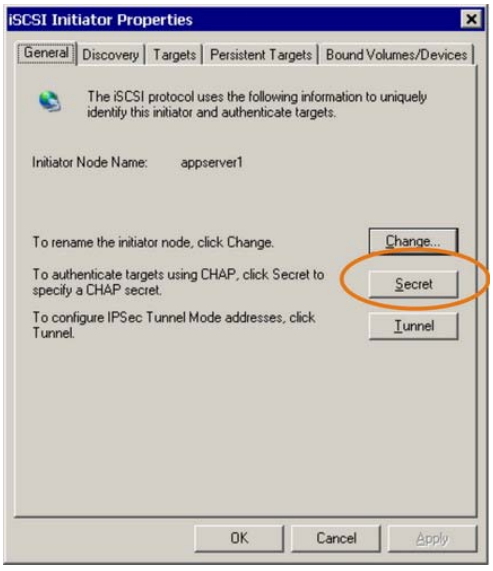

CHAP Secret Setup
The secret allows the initiator to authenticate targets when performing mutual CHAP. Targets must also be configured with this initiator secret.
Enter a secure secret:
12345678901234
Reset
OK
Cancel

2. Under the Discovery tab, click the Add button and enter the IP addresses of the iSCSI ports of your storage system.

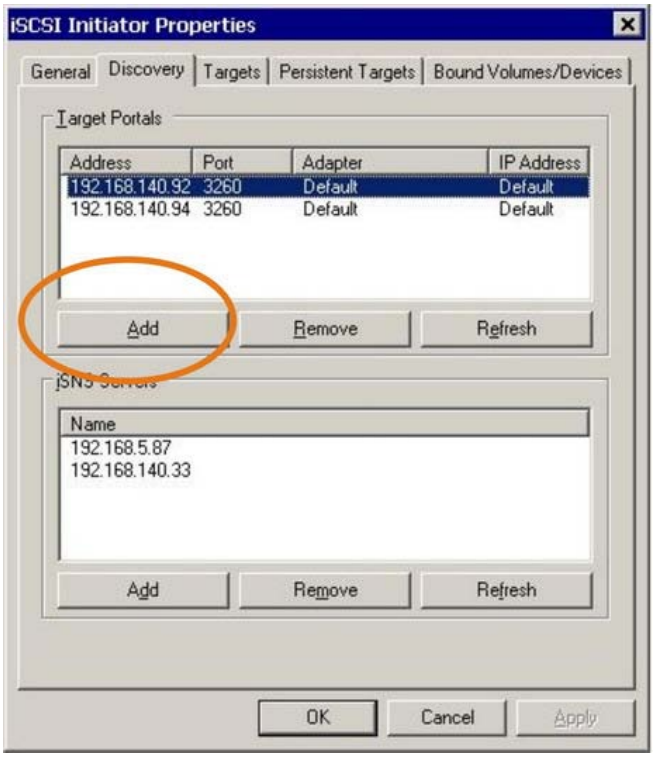

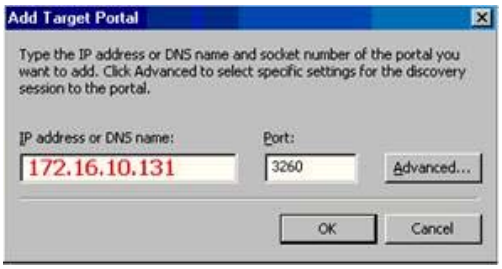
Add Target Portal
Type the IP address or DNS name and socket number of the portal you want to add. Click Advanced to select specific settings for the discovery session to the portal.
IP address or DNS name:
172.16.10.131
Port:
3260
Advanced...
OK
Cancel

3. Under the Targets tab, click Log on to activate the connection with an iSCSI storage target. You should then click the Advanced button after selecting the “Automatically Restore this connection…“ check box. Note that the following procedure is based on the scenario that the iSCSI target has not been logged on.

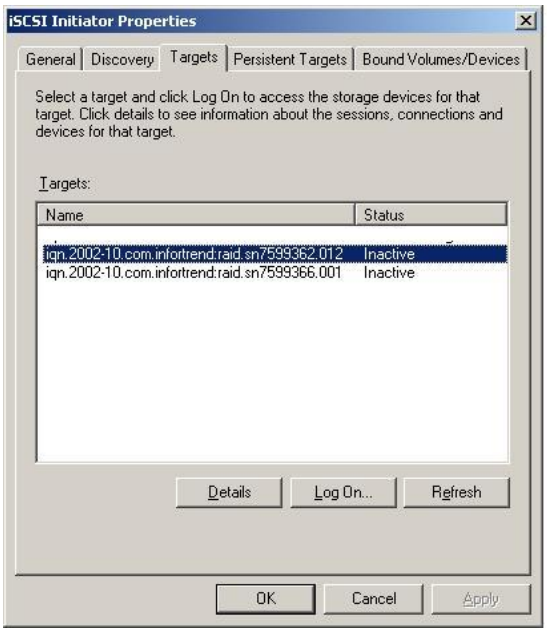

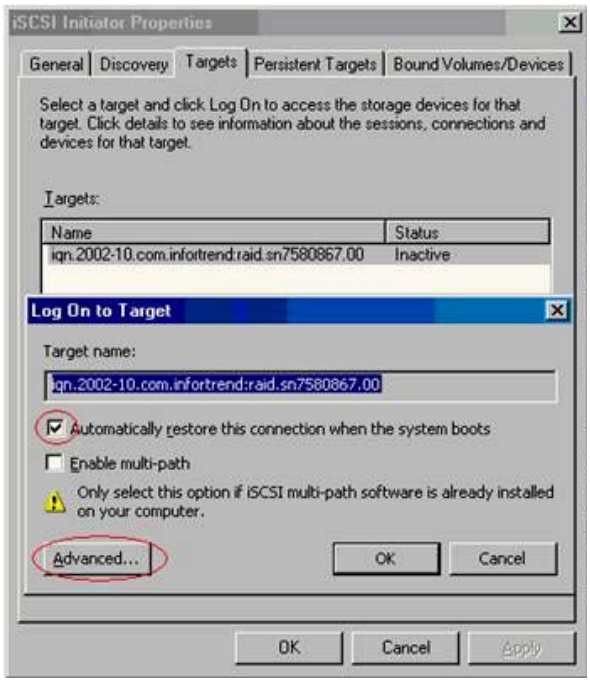

4. Enter User name and Target secret for the One-way CHAP. Click both the CHAP Logon information and the Perform Mutual Authentication check boxes if your prefer using Mutual (Two-way) authentication.

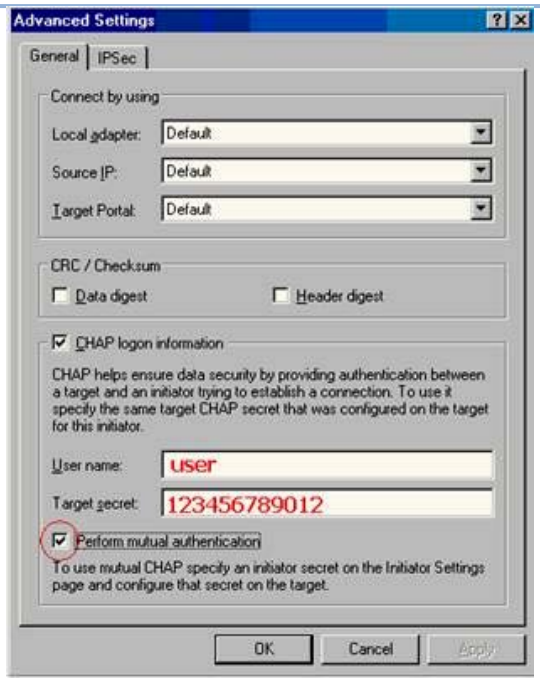
Advanced Settings
General
IPSec
Connect by using
Local adapter:
Default
Source IP:
Default
Target Portal:
Default
CRC / Checksum
Data digest
Header digest
CHAP logon information
CHAP helps ensure data security by providing authentication between a target and an initiator trying to establish a connection. To use it specify the same target CHAP secret that was configured on the target for this initiator.
User name:
user
Target secret:
123456789012
Perform mutual authentication
To use mutual CHAP specify an initiator secret on the Initiator Settings page and configure that secret on the target.
OK
Cancel
Apply

5. You should have correctly configured the Target Portal configuration as described in the previous section. Select appropriate IP and Target Portal settings to designate a host port's relation with an iSCSI target.

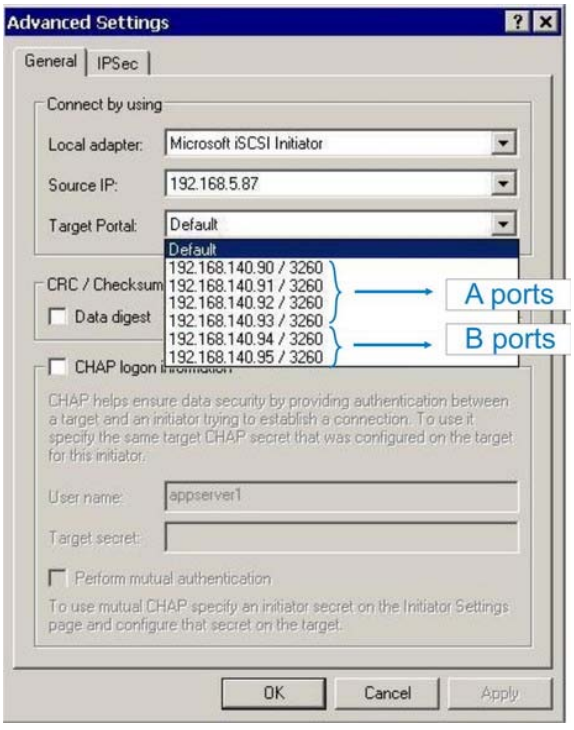

6. Verify a successful connection under the Target tab window. After a short delay, the Target Connection status should be indicated as Connected.

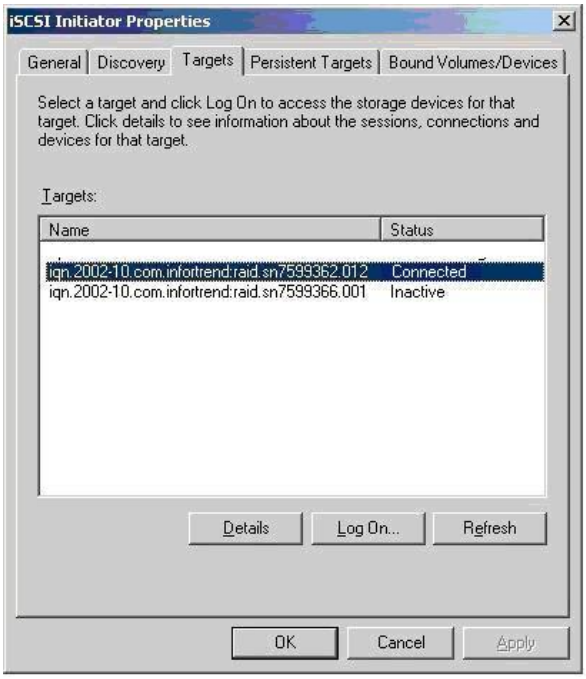

Found New Hardware Wizard
Welcome to the Found New Hardware Wizard
Windows will search for current and updated software by looking on your computer, on the hardware installation CD, or on the Windows Update Web site (with your permission).
Read our privacy policy
Can Windows connect to Windows Update to search for software?
Yes, this time only
Yes, now and every time I connect a device
No, not this time
Click Next to continue.
< Back
Next >
Cancel

You will also be prompted for the Found New Hardware event. Click Cancel to close the message. Galaxy's storage system does not require a device driver.

7. The initiator can have many access routes to a storage target. (Such as a logical drive that appears on 4 host ports). You have to repeat the Add process by clicking the Details button until all related ports are associated with a storage target (a logical drive that appears as an iSCSI target), and therefore the CHAP User Name and Secret have to be entered multiple times for every target portal.

NOTE

Shown below is the corresponding terms used on Galaxy's firmware screen and on Microsoft's initiator. If you see Connected status from the initiator's Discovery screen, then the connection is successful. It is not recommended to change the IQN Node

name on the initiator.

Microsoft iSCSI initiator uses IQN as the default User name for CHAP setting.

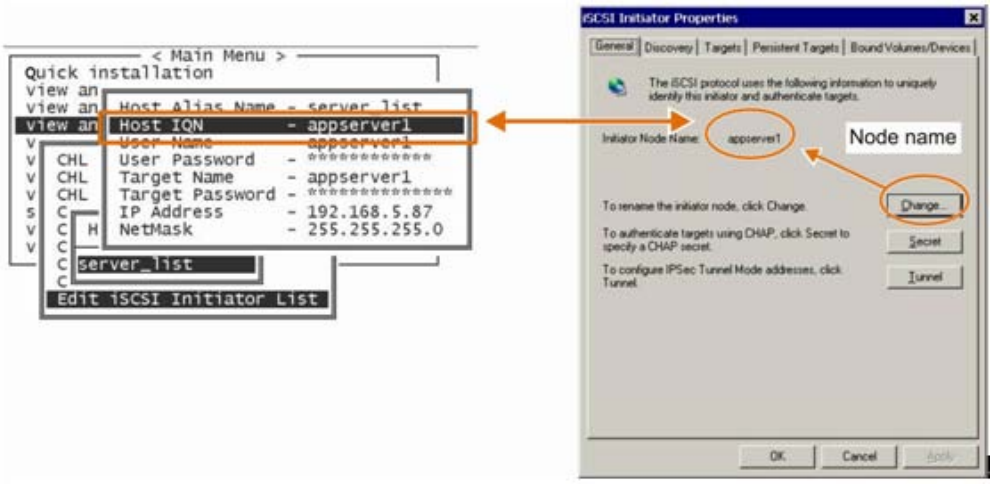

9.7 iSNS Configuration (optional)

9.7.1 iSNS Overview

iSNS stands for Internet Storage Name Service. iSNS is a common discovery, naming, and resource management service for all of the IP storage protocols. Galaxy's iSNS implementation complies with RFC 4171 standards. iSNS discovers iSCSI initiators and targets within a domain and their related information. Windows iSNS server is available in Windows 2000 service pack 4 and Windows Server 2003.

- The iSNS functions can be embedded in an IP Storage switch, gateway, or router, or centralized in an iSNS server.
- Initiators then can query the iSNS to identify potential targets.
- An example of iSNS implementations is Microsoft's iSNS Server 3.0, which is available at Microsoft's download site. http://www.microsoft.com/downloads/details.aspx?familyid=0dbc4af5-9410-4080-a545-f90b45650e20&displaylang=en The iSNS server enables the interchange of data in a domain consisting of initiators and targets according to user's preferences.

9.7.2 iSNS Configuration Sample and Flowchart

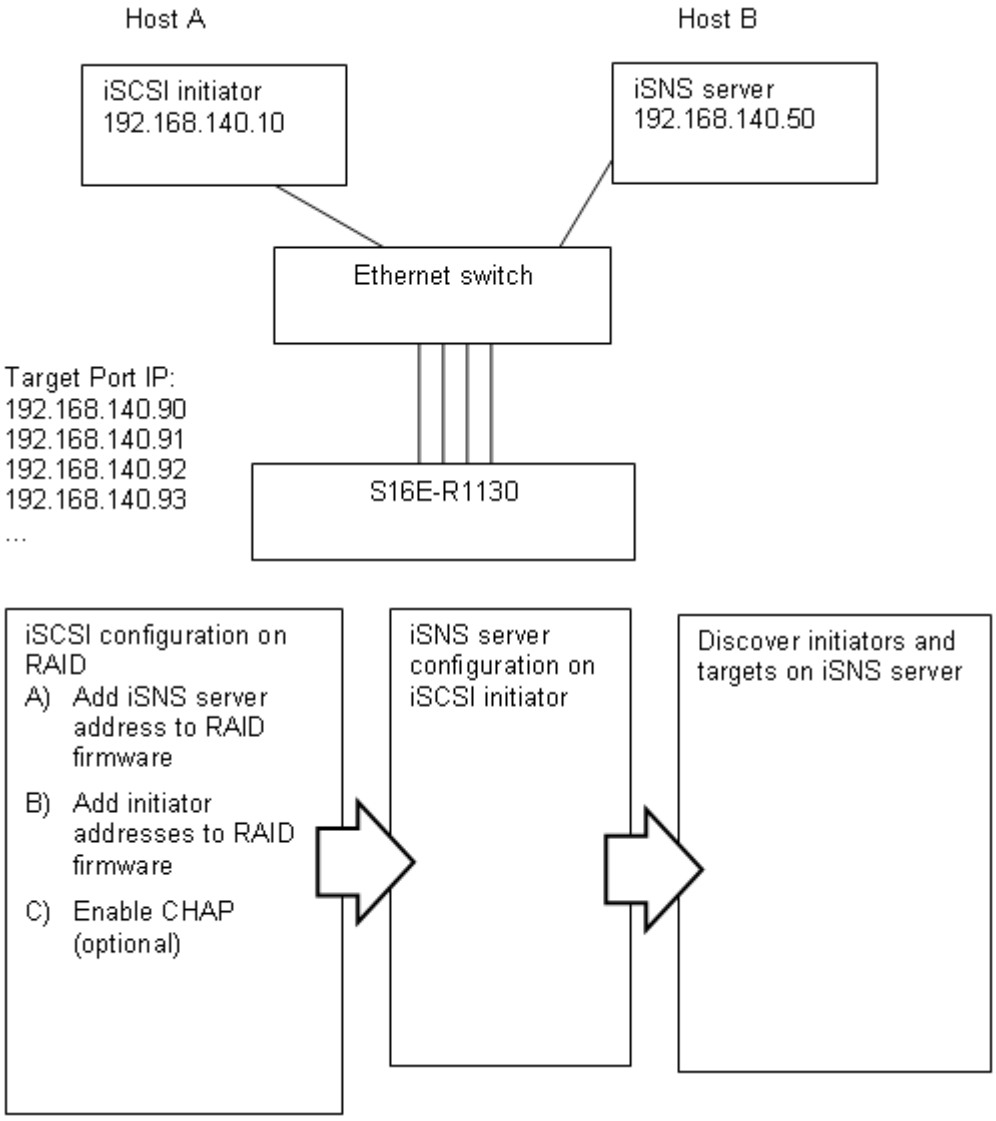

9.7.3 iSNS Configuration (RAID)

Go to: View and Edit Configuration Parameters > Communication Parameters > ISNS Server List

1. Locate the ISNS Server List option in the Communication Parameters window. Press Enter on the IP list screen.

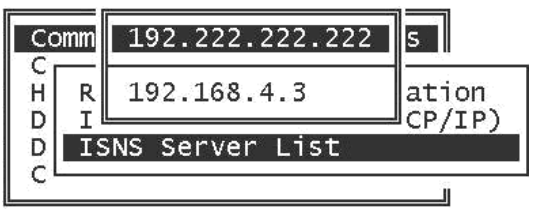

2. Select Add new ISNS server IP Address.

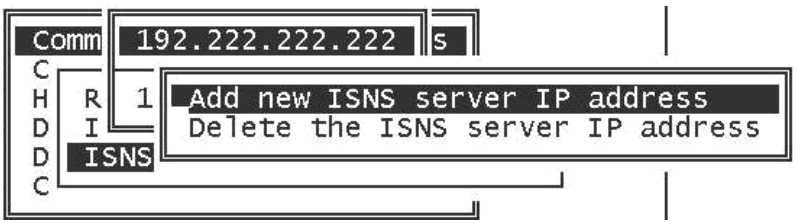

3. Select Yes and then enter the iSNS server address.

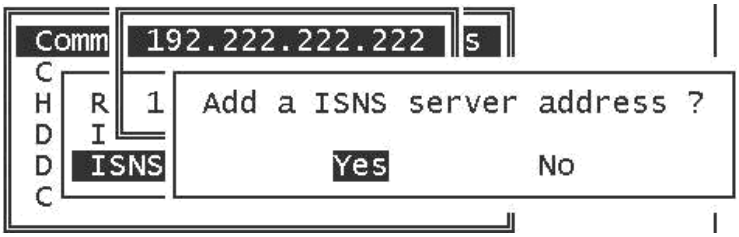

4. Configuring iSNS server address requires resetting the system. You can reset later when you finish configuring other iSCSI parameters.

Please refer to other sections in this chapter for how to configure initiator and CHAP related settings:

9.7.4 iSNS Configuration (PC)

The sample process is based on Microsoft's iSCSI initiator software.

1. Open the iSCSI initiator software, locate the iSNS server field by clicking the Discovery tab.

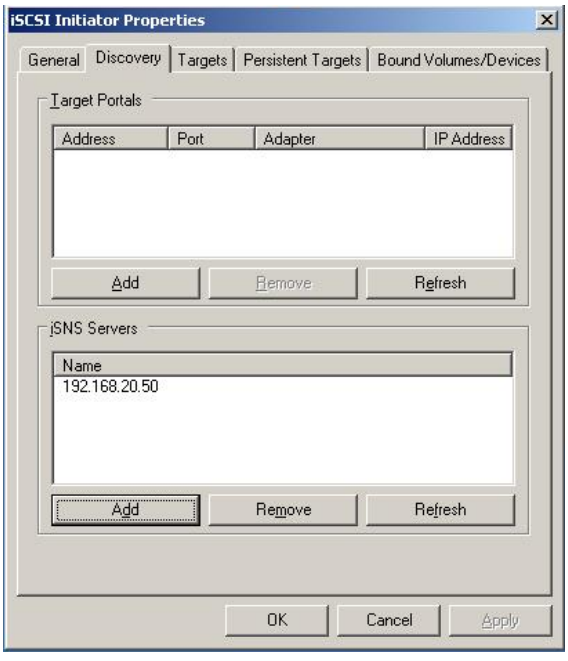

2. Click the Add button to key in an address. After an iSNS server address is added,

you can check on host B (where the iSNS server is installed). And if you have previously configured logical drives and mapped them to host IDs, the target LDs should have been scanned in and appear on the iSNS server configuration screen. Note that an iSNS server may take several minutes to discover devices on the network on the initial setup.

3. If targets and initiators do not show up, please try the Refresh bottom.

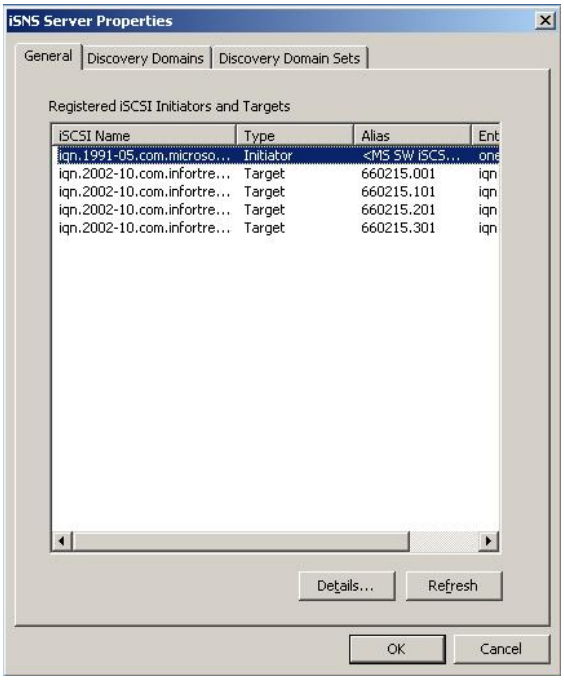

NOTE

An iSNS server is installed and operated using the administrator privilege. An incorrectly installed iSNS can still run, yet the discovery function will not avail.

9.8 SLP Configuration (Optional)

SLP is short for Service Location Protocol, and is supported by the HDX4 firmware. SLP is apt for the following:

- If initiators do not have any information about the target.
- Initiators can either multicast discovery messages directly to the targets or can send discovery messages to storage name servers.

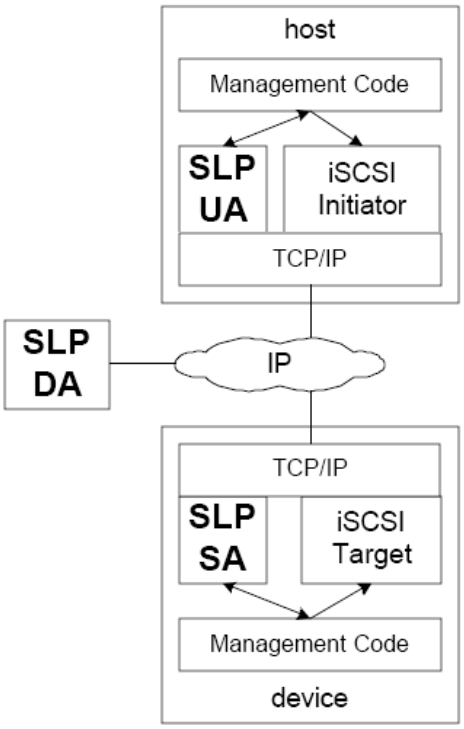

SLP Glossary

- Service Agent (SA): Advertises services, Services have attributes
- User Agent (UA): Finds services, Zero configuration
- Directory Agent (DA), Optional, Propagate service adverts
- SLP Protocol, UDP (default) or TCP, Minimize multicast

9.8.1 How does it work?

1. At startup, UAs and SAs first determine whether there are any DAs on the network.

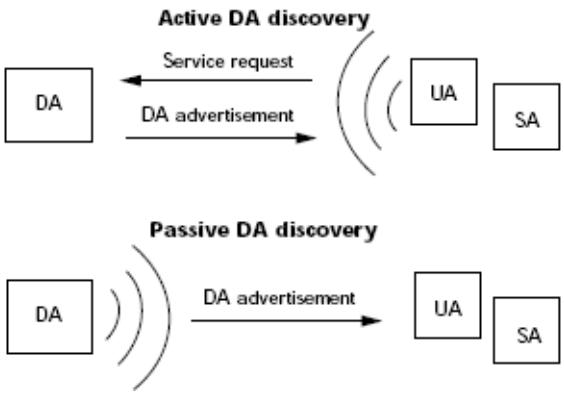

2. A DA is present, it collects all service information advertised by SAs, and UAs

unicast their requests to the DA.

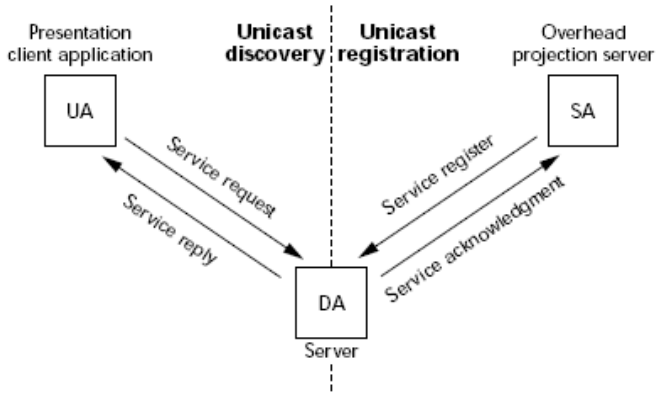

3. In the absence of a DA, UAs repeatedly multicast the same request they would have unicast to a DA. SAs listen for these multicast requests and unicast responses to the UA if it has advertised the requested service.

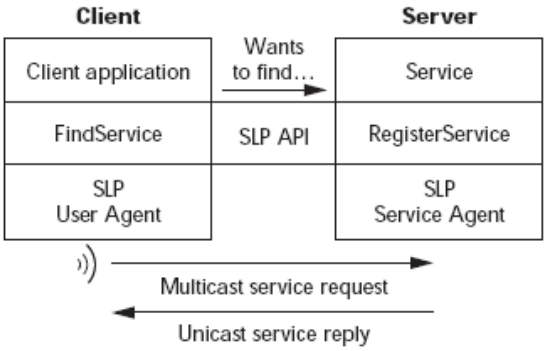

4. The SA registers the service's location with the DA, and the UA obtains this location from the DA in a Service Reply message.

5. Service registrations have lifetimes no greater than 18 hours, so the SA must reregister the service periodically, or the lifetime expires.

9.8.2 Configuring SLP

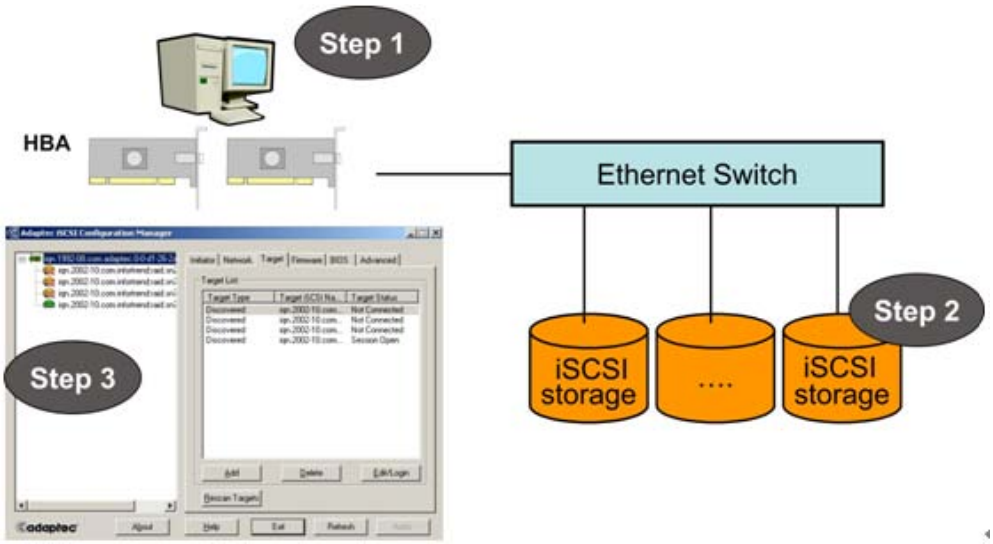

1. Install iSCSI HBAs that support SLP, e.g., an Adaptec TOE.
2. Map logical drives to host
3. Apply LUN Masking if security is a concern.
4. Add your initiators to your firmware initiator list.

Go to: View and Edit Host LUNs > Edit iSCSI Initiator List > (Initiator)

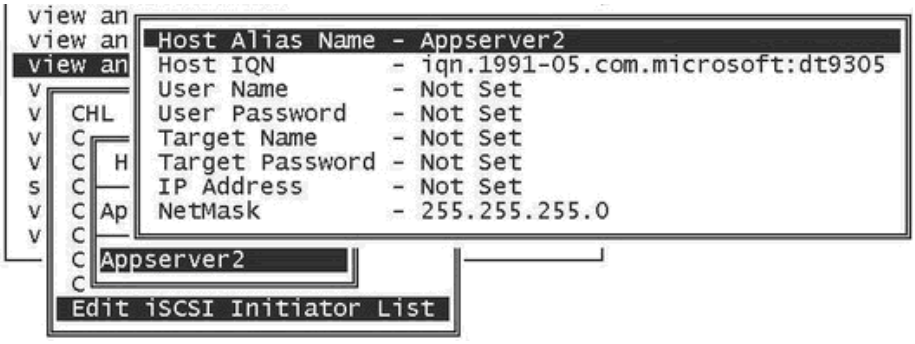
Host Alias Name - Appserver2
Host IQN - iqn.1991-05.com.microsoft:dt9305
User Name - Not Set
User Password - Not Set
Target Name - Not Set
Target Password - Not Set
IP Address - Not Set
NetMask - 255.255.255.0
Appserver2
Edit iSCSI Initiator List

5. After a successful installation, the HBA configuration software should be available on the OS.

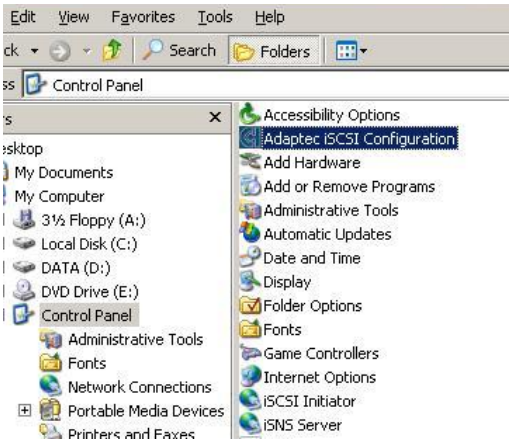

6. Enable the SLP service.

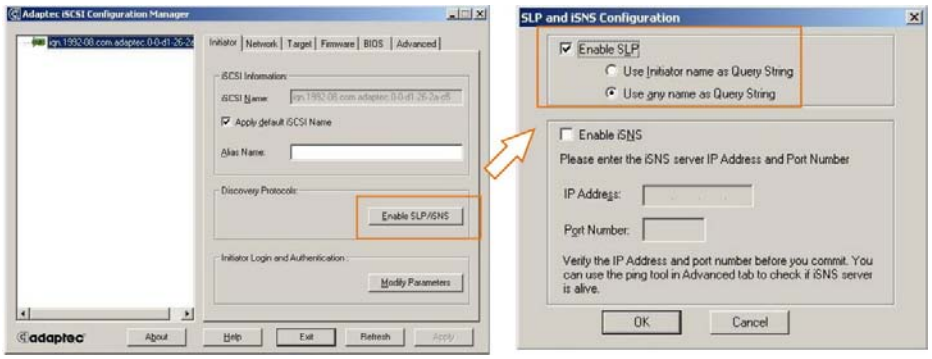
Adaptec iSCSI Configuration Manager
Initiator Network Target Firmware BIOS Advanced
Enable SLP/iSNS
Modify Parameters
About Help Exit Refresh Apply
SLP and iSNS Configuration
Enable SLP
Use Initiator name as Query String
Use any name as Query String
Enable iSNS
Please enter the iSNS server IP Address and Port Number
IP Address:
Port Number:
Verify the IP Address and port number before you commit. You can use the ping tool in Advanced tab to check if iSNS server is alive.
OK
Cancel

7. Under the Target tab, iSCSI targets should have been scanned in after a while, or you may use Rescan Targets button.

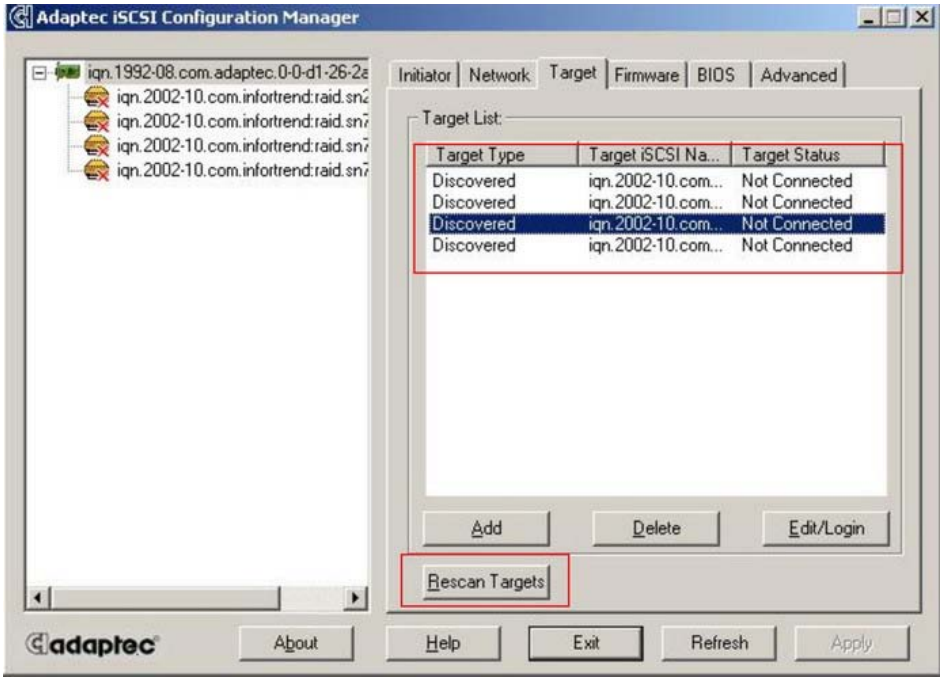

8. Activate the access to targets by logging in, and apply CHAP authentication if

preferred.

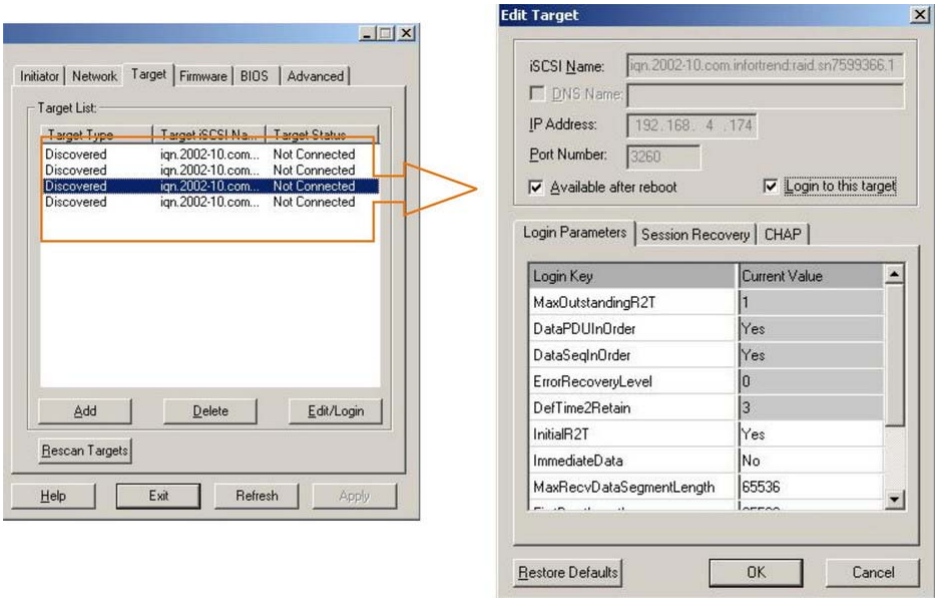

9. Check the iSCSI session status and complete the disk drive initialization process.

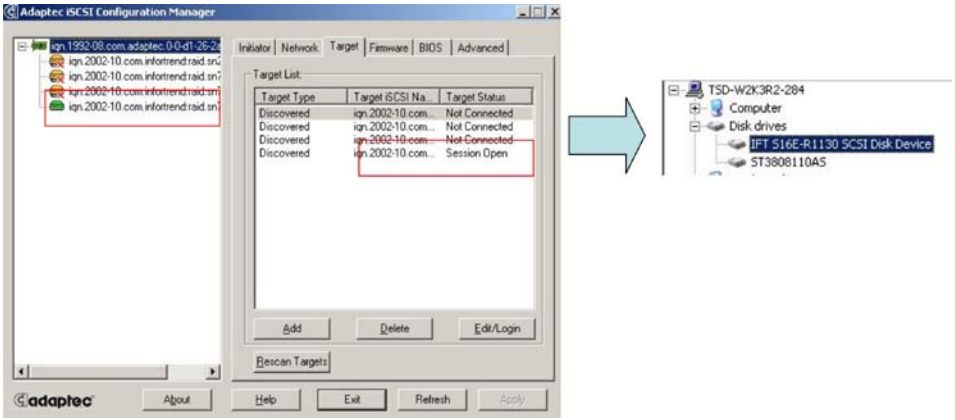

9.9 Jumbo Frame

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Jumbo Frames

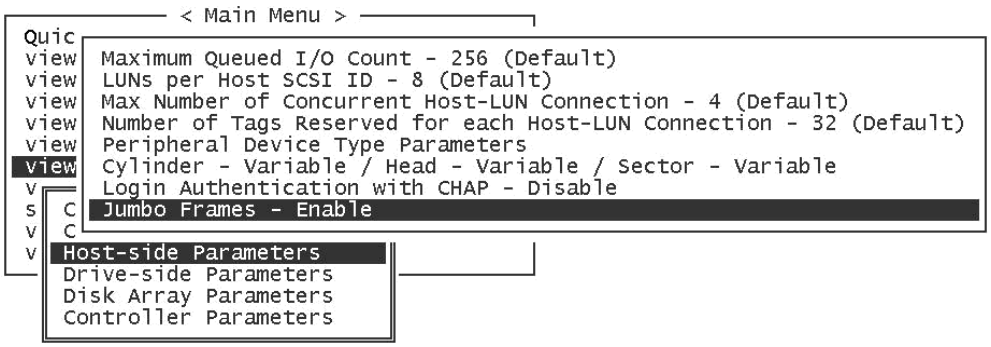
< Main Menu >
Maximum Queued I/O Count - 256 (Default)
LUNs per Host SCSI ID - 8 (Default)
Max Number of Concurrent Host-LUN Connection - 4 (Default)
Number of Tags Reserved for each Host-LUN Connection - 32 (Default)
Peripheral Device Type Parameters
Cylinder - Variable / Head - Variable / Sector - Variable
Login Authentication with CHAP - Disable
Jumbo Frames - Enable
Host-side Parameters
Drive-side Parameters
Disk Array Parameters
Controller Parameters

Jumbo Frames extend Ethernet's bytes per frame size, and can significantly increase performance.

NOTE

The Jumbo Frame feature requires that all of the end devices in an iSCSI network support and have their Jumbo Frame function activated. Also pay attention to the maximum frame sizes set with Jumbo Frames on these devices.

10 Host-side and Drive-side Parameters

This chapter discusses the advanced options for tuning various firmware parameters. Each function is given a brief explanation as well as a configuration sample. Terminal screens are used in the configuration samples. Some of the operations require basic knowledge of RAID technology and are only recommended for an experienced user.

NOTE

All figures in this chapter are showing examples using the management hyper terminal screen.

10.1 Host-side Parameters

he controller supports the following Host-side configurations:

- Maximum Queued I/O Count
- LUNs per Host ID
- Max. Number of Concurrent Host-LUN Connection
- Tags per Host-LUN Connect
- Peripheral Dev Type Parameters
- Cyl/Head/Sector Mapping Config

10.1.1 Maximum Concurrent Host LUN Connection (“Nexus” in SCSI)

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Max Number of Concurrent Host-Lun Connection

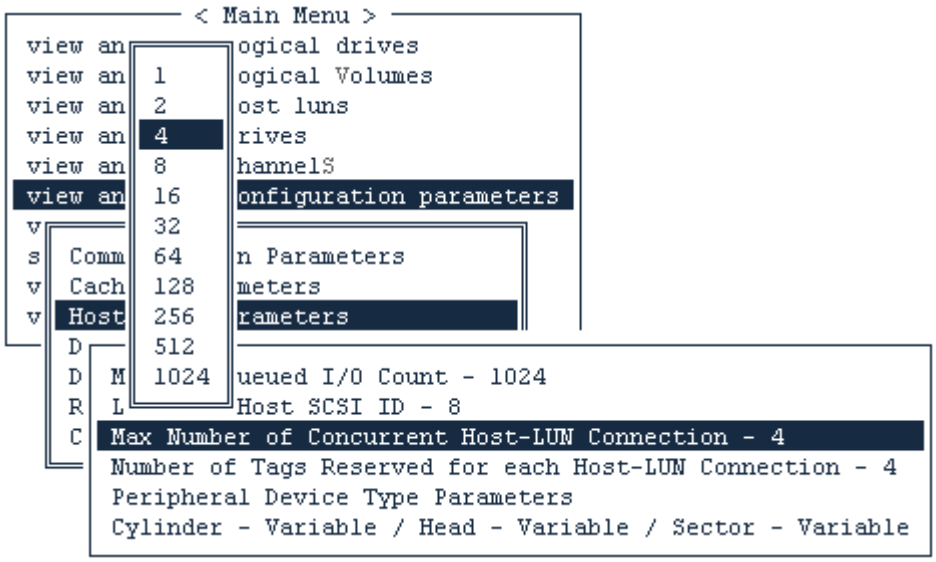

The "Max Number of Concurrent Host-LUN Connection" menu option is used to set the maximum number of concurrent host-LUN connections. Change this menu option setting only if you have more than four logical drives or partitions. Increasing this number might increase your performance.

Maximum concurrent host LUN connection (nexus in SCSI) is the arrangement of the controller internal resources for use with a number of the current host nexus.

For example, you can have four hosts (A, B, C, and D) and four host IDs/LUNs (IDs 0, 1, 2 and 3) in a configuration where:

- Host A accesses ID 0 (one nexus).
- Host B accesses ID 1 (one nexus).
- Host C accesses ID 2 (one nexus).
- Host D accesses ID 3 (one nexus).

These connections are all queued in the cache and are called four nexus.

If there is I/O in the cache with four different nexus, and another host I/O comes with a nexus different than the four in the cache (for example, host A accesses ID 3), the controller returns busy. This occurs with the concurrent active nexus; if the cache is cleared, it accepts four different nexus again. Many I/O operations can be accessed via the same nexus.

NOTE

The maximum number of tags or concurrent Host-LUN connections is determined by

the following:

LUN-per-ID x tags reserved = flag A

Max. Number of Concurrent Host-LUN connection = flag B

If A>B, Max.=A; else, Max.=B

10.1.2 Number of Tags Reserved for Each Host-LUN Connection

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Number of Tags Reserved for Each Host-LUN Connection

Each nexus has 32 (the default setting) tags reserved. When the host computer sends 8 I/O tags to the controller, and the controller is too busy to process them all, the host might start to send less than 8 tags during every certain period of time since then. This setting ensures that the controller will accept at least 32 tags per nexus.

The controller will be able to accept more than that as long as the controller internal resources allow - if the controller does not have enough resources, at least 32 tags can be accepted per nexus.

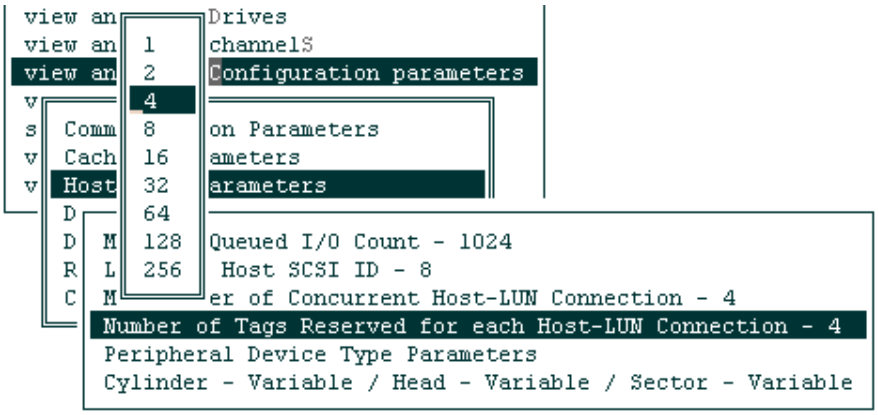

10.1.3 Maximum Queued I/O Count

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Maximum Queued I/O Count

The "Maximum Queued I/O Count" menu option enables you to configure the maximum number of I/O operations per host channel that can be accepted from servers. The predefined range is from 1 to 1024 I/O operations per host channel, or you can choose the "Auto" (automatically configured) setting. The default value is 256 I/O operations. The maximum number of queued I/O operations is 4096.

The appropriate "Maximum Queued I/O Count" setting depends on how many I/O operations attached servers are performing. This can vary according to the amount of host memory present as well as the number of drives and their size. If you increase the amount of host memory, add more drives, or replace drives with higher performance, you might want to increase the maximum I/O count. But usually optimum performance results from using the "Auto" or "256" settings.

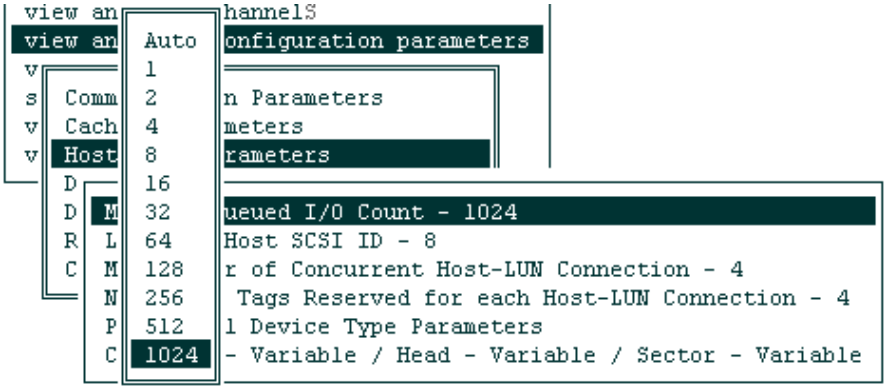

10.1.4 LUNs per Host ID

Go to: View and Edit Configuration Parameters > Host-Side Parameters > LUNs per Host ID

The highly scalable Fibre Channel technology can address up to 126 devices per loop, and theoretically more than a million using FC switches. Each configured RAID volume is associated with host IDs and appears to the host as a contiguous volume.

If you file a document into a cabinet, you must put the document into one of the drawers. As defined by storage interface architecture, a Fibre channel ID is like a cabinet, and the drawers are the LUNs (Logical Unit Numbers). Each Fibre channel ID encapsulates up to 32 LUNs and up to 1024 LUNs are configurable through all host ports. A RAID volume can be associated with any of the LUNs under the Fibre channel IDs. Most Fibre host adapters treats a LUN like another Fibre device.

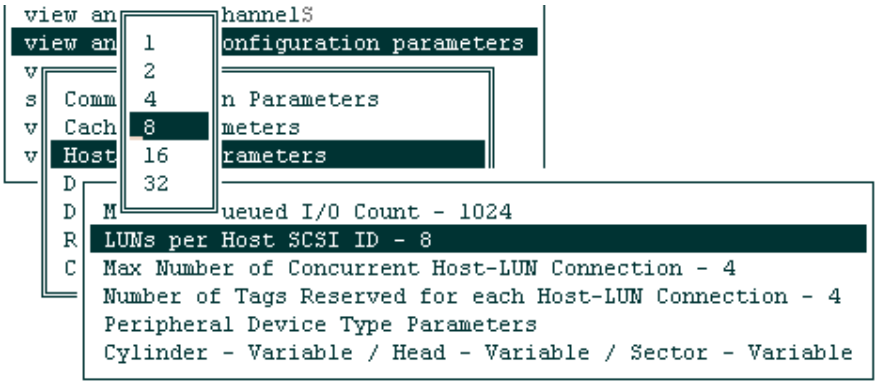
1
2
4
8
16
32
ueued I/O Count - 1024
LUNs per Host SCSI ID - 8
Max Number of Concurrent Host-LUN Connection - 4
Number of Tags Reserved for each Host-LUN Connection - 4
Peripheral Device Type Parameters
Cylinder - Variable / Head - Variable / Sector - Variable

10.1.5 LUN Applicability

The LUN Applicability settings apply in environments where system administrators use in-band methodology for management access to a RAID subsystem.

If no logical drive has been created and mapped to a host LUN, and the RAID controller is the only device connected to the host computer, usually the operating system will not load the driver for the host adapter. If the driver is not loaded, the host computer will not be able to use the in-band utility to communicate with the RAID controller. This is often the case when users want to start configuring a brand new subsystem using Galaxy Array Manager software.

The Galaxy HDX4 series firmware automatically assigns the “First Undefined LUN” to enable in-band access. Even if a mapped volume becomes unmapped, in-band management is still valid.

10.1.6 Peripheral Device Type

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Peripheral Device Type Parameters > Peripheral Device Type

Firmware default is Enclosure Service Device, which enables a brand new system to appear to host to enable in-band management. You do not have to change the default.

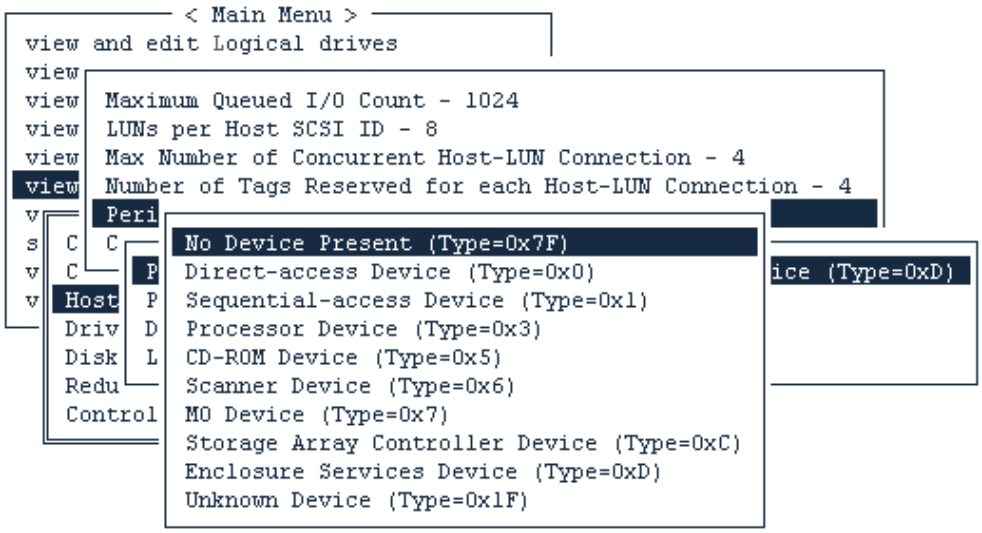
< Main Menu >
view and edit Logical drives
Maximum Queued I/O Count - 1024
LUNs per Host SCSI ID - 8
Max Number of Concurrent Host-LUN Connection - 4
Number of Tags Reserved for each Host-LUN Connection - 4
No Device Present (Type=0x7F)
Direct-access Device (Type=0x0)
Sequential-access Device (Type=0x1)
Processor Device (Type=0x3)
CD-ROM Device (Type=0x5)
Scanner Device (Type=0x6)
MO Device (Type=0x7)
Storage Array Controller Device (Type=0xC)
Enclosure Services Device (Type=0xD)
Unknown Device (Type=0x1F)

10.1.7 In-band Management Access

External devices (including a RAID subsystem; from the view of an application server or management PC) require communication links with a management computer for

device monitoring and administration. In addition to the regular RS-232C or Ethernet connection, in-band SCSI can serve as an alternative means of management communications. In-band SCSI translates the original configuration commands into standard SCSI commands. These SCSI commands are then sent to and received by the controller over the existing host links, either SCSI or Fibre.

10.1.8 Peripheral Device Type Parameters for Various Operating Systems

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Peripheral Device Type Parameters

NOTE

There is no need to configure the Peripheral Device setting if you are trying to manage a RAID subsystem from a Galaxy Array Manager station through an Ethernet connection (to the Galaxy HDX4 subsystem's Ethernet port). An Ethernet connection to RAID uses TCP/IP as the communication protocol.

Do not change the firmware peripheral device default settings.

Different host operating systems require different adjustments. See the tables below to find appropriate settings for your host operating system. References to "Peripheral Device Qualifier" and "Device Support for Removable Media" are also included.

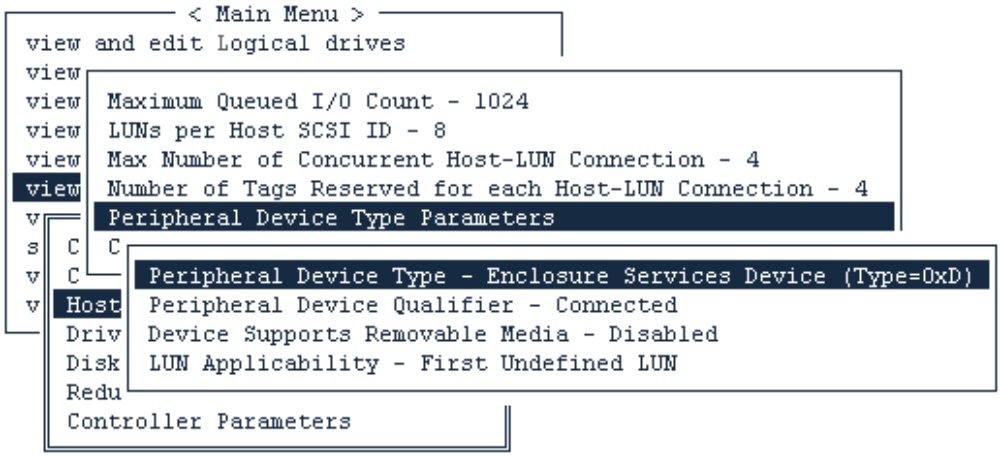
< Main Menu >
view and edit Logical drives
Maximum Queued I/O Count - 1024
LUNs per Host SCSI ID - 8
Max Number of Concurrent Host-LUN Connection - 4
Number of Tags Reserved for each Host-LUN Connection - 4
Peripheral Device Type Parameters
Peripheral Device Type - Enclosure Services Device (Type=0xD)
Peripheral Device Qualifier - Connected
Device Supports Removable Media - Disabled
LUN Applicability - First Undefined LUN
Controller Parameters

Operating System	Peripheral Device Type	Peripheral Device Qualifier	Device Support for Removable Media	LUN Applicability
Windows	0xd	Connected	Either is okay	First Defined

2000/2003				LUN
Solaris™ 8/9 (x86 and SPARC)	0xd	Connected	Either is okay	First Defined LUN
Linux RedHat 8/9; SuSE 8/9	0xd	Connected	Either is okay	First Defined LUN

Peripheral Device Type Settings

Device Type	Setting
Enclosure Service Device	0xd
No Device Present	0x7f
Direct-access Device	0
Sequential-access Device	1
Processor Type	3
CD-ROM Device	5
Scanner Device	6
MO Device	7
Storage Array Controller Device	0xC
Unknown Device	0x1f

10.1.9 Cylinder/Head/Sector Mapping

Go to: View and Edit Configuration Parameters > Host-Side Parameters > Cylinder

Drive capacity is decided by the number of blocks. For some operating systems (Sun Solaris, for example) the capacity of a drive is determined by the cylinder/head/sector count. For earlier Sun Solaris systems, the cylinder cannot exceed 65535; choose "cylinder<65535," then the controller will automatically adjust the head/sector count for your OS to read the correct drive capacity. Please refer to the related documents provided with your operating system for more information.

Cylinder, Head, and Sector counts are selectable from the configuration menus shown below. To avoid any difficulties with a Sun Solaris configuration, the values

listed below can be applied.

Capacity	Cylinder	Head	Sector
< 64 GB	variable	64	32
64 - 128 GB	variable	64	64
128 – 256 GB	variable	127	64
256 – 512 GB	variable	127	127
512 GB - 1 TB	variable	255	127

Older Solaris versions do not support drive capacities larger than 1 terabyte. Solaris 10 now supports array capacity larger than 1TB. Set the values to the values listed in the table below:

Capacity	Cylinder	Head	Sector
>1TB	<65536	255	variable
		variable	255

Configuring Sector Ranges/Head Ranges/Cylinder Ranges:

The sector, head, and cylinder variables are presented as preset combinations. Please refer to the documentation that came with your operating system and select one value set that is most appropriate for your OS file system.

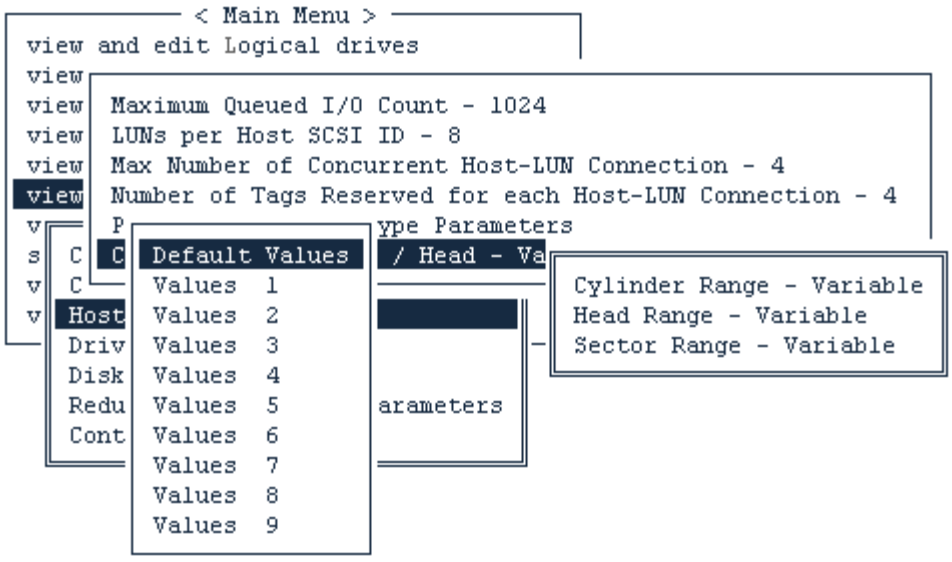

10.2 Drive-side Parameters

Go to: View and Edit Configuration Parameters > Drive-Side Parameters

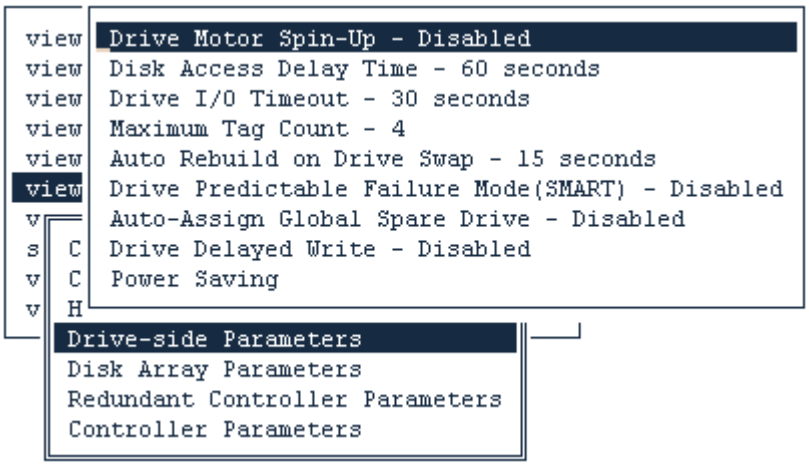
Drive Motor Spin-Up - Disabled
Disk Access Delay Time - 60 seconds
Drive I/O Timeout - 30 seconds
Maximum Tag Count - 4
Auto Rebuild on Drive Swap - 15 seconds
Drive Predictable Failure Mode(SMART) - Disabled
Auto-Assign Global Spare Drive - Disabled
Drive Delayed Write - Disabled
Power Saving
Drive-side Parameters
Disk Array Parameters
Redundant Controller Parameters
Controller Parameters

10.2.1 Disk Access Delay Time

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Disk Access Delay Time

This feature sets the delay time before the subsystem tries to access the hard drives after power-on. Default may vary from 15 seconds to 30 seconds, and is determined by the type of drive interface. This parameter can be adjusted to fit the spin-up speed of different disk drive models.

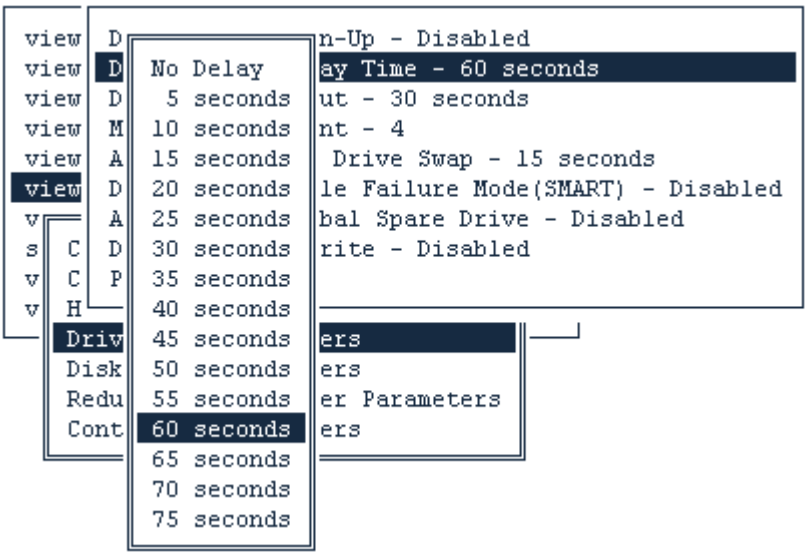

10.2.2 Drive I/O Timeout

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Drive I/O Timeout

The “Drive I/O Timeout” is the time interval for the controller to wait for a drive to respond. If the controller attempts to read data from or write data to a drive but the drive does not respond within the Drive I/O Timeout value, the drive will be considered as a failed drive.

When the drive itself detects a media error while reading from the drive platter, it usually retries the previous reading or re-calibrates the read/write head. When a disk drive encounters a bad block on the media, it will attempt to reassign the bad block to a spare block. However, it takes time to perform the above operations. The time to perform these operations can vary between among disk drives by different vendors.

During channel bus arbitration, a device with higher priority can utilize the bus first. A device with lower priority will sometimes receive an I/O timeout when devices of higher priority keep utilizing the bus.

The default setting for “Drive I/O Timeout” is 7 seconds. It is highly recommended not to change this setting. Setting the timeout to a lower value will cause the controller to judge a drive as failed while a drive is still retrying, or while a drive is unable to arbitrate the drive bus. Setting the timeout to a greater value will cause the controller to keep waiting for a drive, and it may sometimes cause a host timeout.

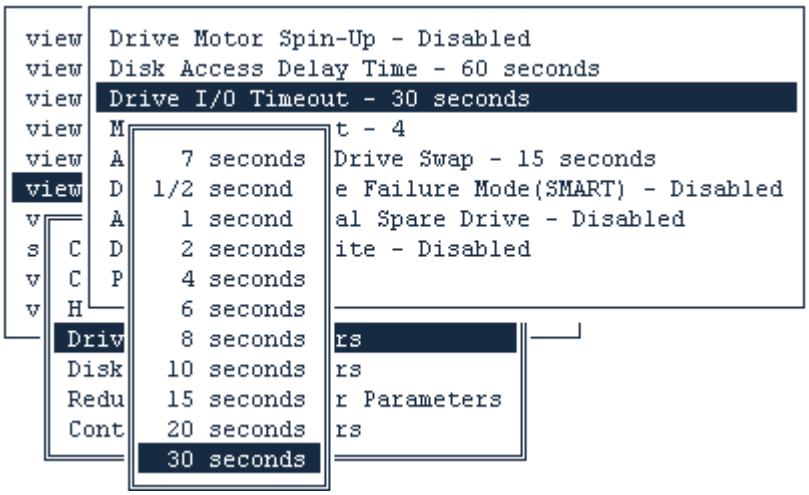

10.2.3 Maximum Tag Count: Tag Command Queuing (TCQ) and Native Command Queuing (NCQ) Support

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Maximum

Tag Count

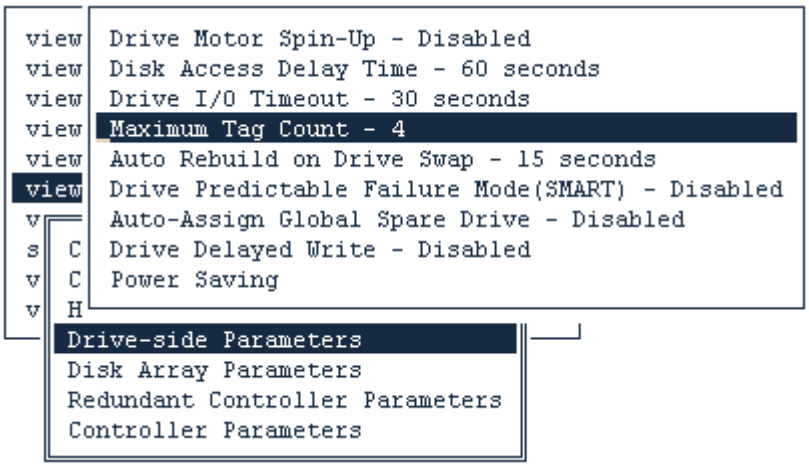

This sub-menu facilitates the support for both Tagged Command Queuing (TCQ) and Native Command Queuing (NCQ). TCQ is a traditional feature on SCSI, SAS, or Fibre Channel disk drives, while NCQ is recently implemented with SATA disk drives. The queuing feature requires the support of both host adapters and hard disk drives. Command queuing can intelligently reorder host requests to streamline random accesses for IOPS/multi-user applications.

Galaxy's subsystems support Tag Command Queuing with an adjustable maximum tag count from 1 to 128. The default setting is "Enabled" with a maximum tag count of 32 (SCSI), 8 (for Fibre drives), or 4 (default for SAS/SATA drives).

NOTE

Every time you change this setting, you must reset the controller/subsystem for the changes to take effect.

Disabling Tag Command Queuing will disable the hard drives' built-in buffer.

The following options are categorized as related to array maintenance and data integrity:
Auto Rebuild on Drive Swap
Auto-Assign Global Spare Drive

Another option is associated with disk drive S.M.A.R.T. support.

10.2.4 Drive Delayed Write

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Drive Delayed Write

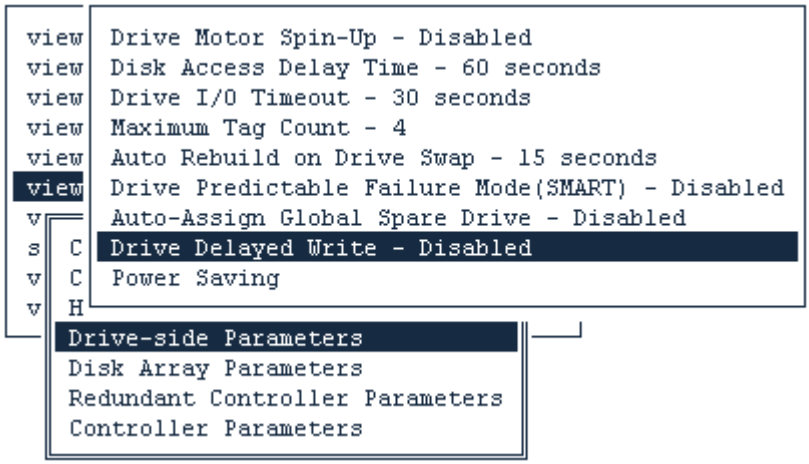

This option applies to disk drives which come with embedded buffers. When enabled, write performance can improve. However, this option should be disabled for mission-critical applications. In the event of power outage or drive failures, data cached in drive buffers will be lost, and data inconsistency will occur.

10.2.5 Power Saving

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Power Saving

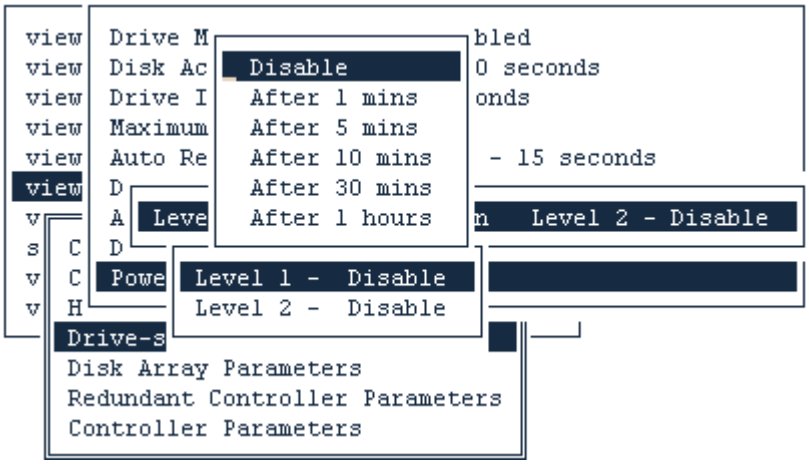

This feature supplements the disk spin-down function, and supports power-saving on specific logical drives or un-used disk disks with an idle state and the 2-stage power-down settings.

Advantages: see the power saving features below.

Applicable Disk Drives: Logical drives and non-member disks [including spare drives and un-used drives (new or formatted drives)]. The power-saving policy set to an individual logical drive (from the View and Edit Logical Drive menu) has priority

over the general Drive-side Parameter setting.

Power-saving Levels:

Level	Power Saving Ratio	Recovery Time	ATA command	SCSI command
Level 1 (Idle) *	19% to 22% **	1 second	Idle	Idle
Level 2 (Spin-down) *	80%	30 to 45 seconds	Standby	Stop

Note

- The Idle and Spin-down modes are defined as Level 1 and Level 2 power saving modes on Galaxy's user interfaces.
- The power saving ratio is deducted by comparing the consumption in idle mode against the consumption when heavily stressed.
- Hard drives can be configured to enter the Level 1 idle state for a configurable period of time before entering the Level 2 spin-down state.
- Four power-saving modes are available:
  Disable,
  Level 1 only,
  Level 1 and then Level 2,
  Level 2 only. (Level 2 is equivalent to legacy spin-down)
- The Factory defaults is "Disabled" for all drives. The default for logical drives is also Disabled.
- The preset waiting period before entering the power-saving state:
  Level 1: 5 minutes
  Level 2: 10 minutes (10 minutes from being in the level 1)
- If a logical drive is physically relocated to another enclosure (drive roaming), all related power-saving feature is cancelled.

Applicable Hardware:

- All Galaxy HDX4 series running the compatible firmware version.

11 Enclosure Management

This chapter discusses the configuration options related to enclosure monitoring. Each function is given a brief explanation as well as a configuration sample. Terminal screens will be used in the configuration samples. Some of the operations require basic knowledge of RAID technology and are only recommended for an experienced user.

NOTE

All figures in this chapter are showing examples from a hyper terminal console.

11.1 Enclosure Device Statuses (Peripheral Device Status)

11.1.1 RAID Enclosure Devices

Go to: View and Edit Peripheral Devices > View Peripheral Device Status > I2C Peripheral Device

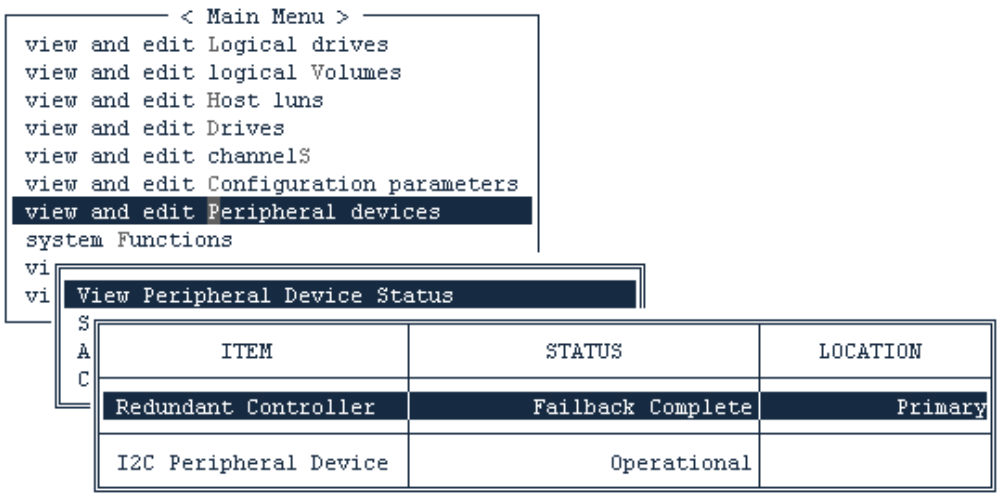

Press [ENTER] on the “SES Device” or “I2C Peripheral Device” to display a list of peripheral devices (enclosure modules).

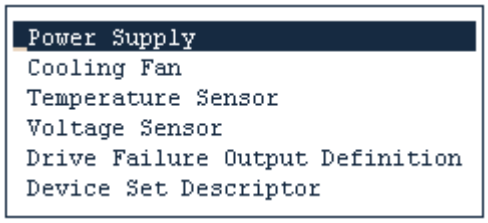

Monitoring of device status depends on enclosure implementation and is accessed through different interfaces, e.g., S.E.S., SAS wide links, or I2C serial bus. Below is a screen showing the enclosure devices interfaced through an I2C serial bus:

NOTE

A SAS expansion enclosure connected through SAS links is also considered as an I2C Peripheral Device, which is defined as the Device Set 1 (JBOD enclosure device) next to the Device Set 0 (RAID enclosure device).

Press [ENTER] on a component type to examine its operating status. Following is a screen listing all cooling fans in a 3U enclosure, including those embedded in power supply modules.

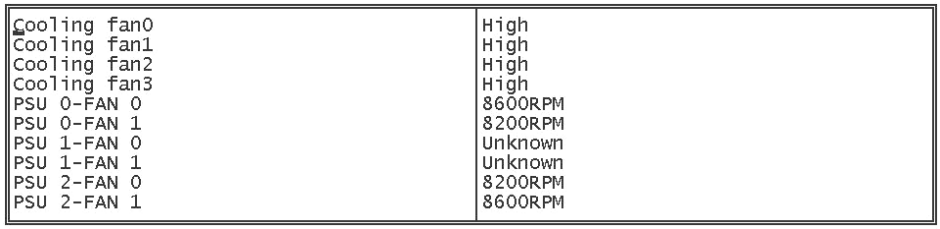

11.1.2 Devices within the Expansion Enclosure

Devices in SAS expansion enclosures are monitored through a proprietary in-band methodology through a monitor chipset on JBOD controllers. Below is the device shown on the View and Edit Drives screen.

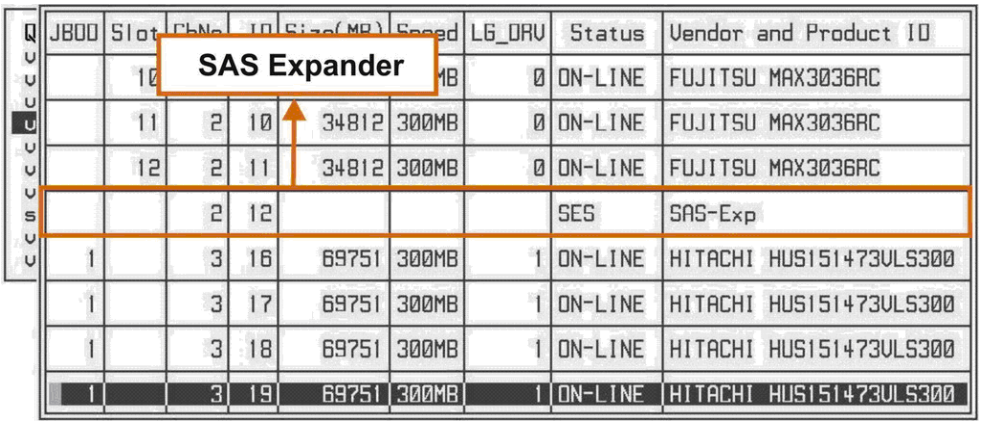
SAS Expander
JBOD Slot ChNo ID Size(MB) Speed LG_DRV Status Vendor and Product ID
10 MB 0 ON-LINE FUJITSU MAX3036RC
1 11 2 10 34812 300MB 0 ON-LINE FUJITSU MAX3036RC
1 12 2 11 34812 300MB 0 ON-LINE FUJITSU MAX3036RC
2 12 SES SAS-Exp
1 3 16 69751 300MB 1 ON-LINE HITACHI HUS151473VLS300
1 3 17 69751 300MB 1 ON-LINE HITACHI HUS151473VLS300
1 3 18 69751 300MB 1 ON-LINE HITACHI HUS151473VLS300
1 3 19 69751 300MB 1 ON-LINE HITACHI HUS151473VLS300

Information about the SAS expander handling SAS expansion links is shown as the last device in the RAID enclosure.

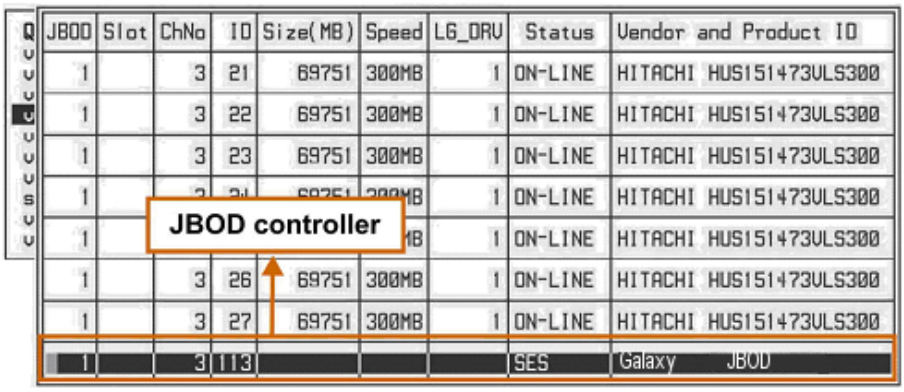

JBOD	Slot	ChNo	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1		3	21	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	22	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	23	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1						1	ON-LINE	HITACHI HUS151473VLS300
1						1	ON-LINE	HITACHI HUS151473VLS300
1		3	26	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	27	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	113				SES	Galaxy JBOD

The JBOD controller within the expansion enclosure is shown as the last device in the expansion enclosure. You may press [ENTER] on the device to check the revision number of the firmware running on SAS channel devices.

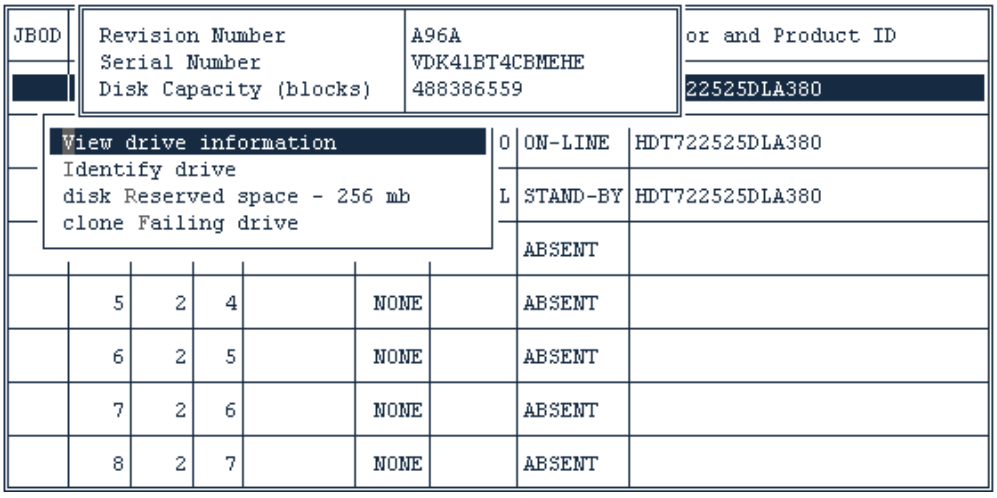

The operating statuses of individual enclosure devices within the expansion enclosures can be found in View and Edit Peripheral Devices > View Peripheral Device Status > I2C Peripheral Device

NOTE

The JBOD enclosure devices will only display when firmware detects expansion enclosures across its expansion links.

11.1.3 Verifying Disk Drive Failure in a Multi-enclosure Application

Go to: View and Edit Peripheral Devices > View Peripheral Device Status > I2C

Peripheral Device > Drive Failure Output Definition

You can verify disk drive locations by checking their channel number, slot number, and device IDs in “Drive Failure Output Definition.” Note that the SAS channel number is a logically defined congregation of multiple physical links (PHYs) through the SAS expanders.

This information is important for locating and replacing a failed drive.

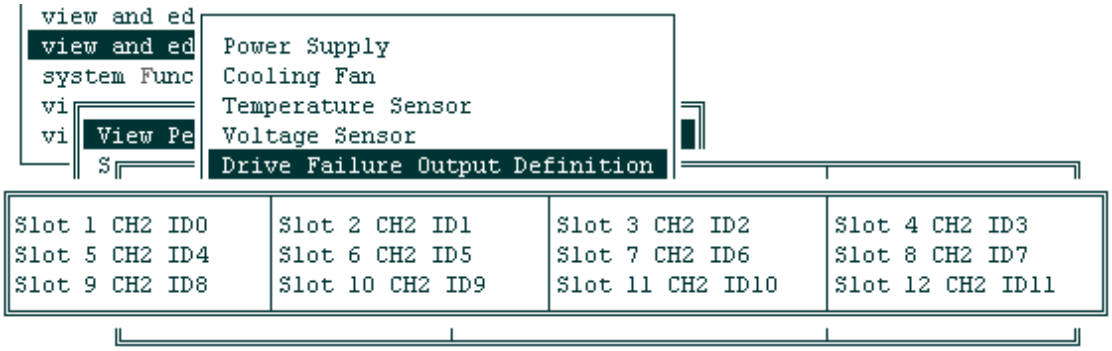

Another key factor in identify drive location is the JBOD/SBOD identifier that can be found under the Main Menu -> “View and Edit Drives” sub-menu. The JBOD identifier equals the enclosure ID you configure using the front panel rotary switch or the rear panel DIP switches.

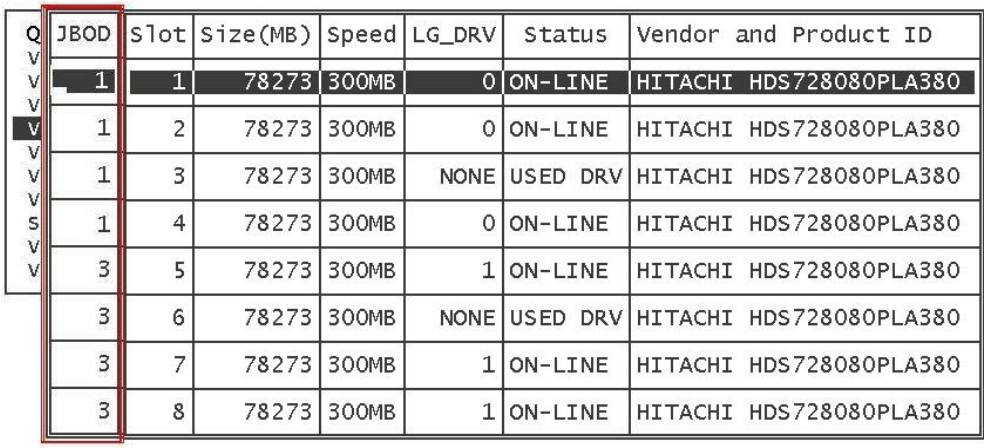

JBOD	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	1	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	2	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	3	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
1	4	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
3	5	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	6	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
3	7	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	8	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380

11.2 Enclosure Management Options

11.2.1 Event Triggered Operations

Go to: View and Edit Peripheral Devices > Set Peripheral Device Entry > Event Trigger Operations

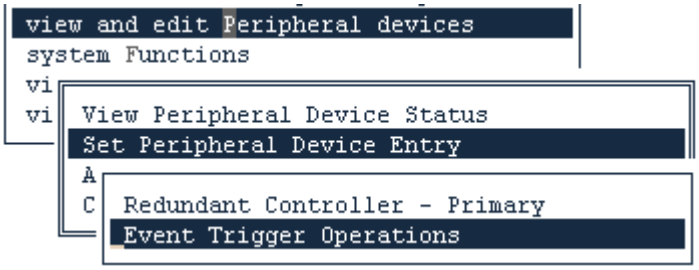

1. Use arrow keys to move your cursor bar to select “View and Edit Peripheral Devices” on the Main Menu and press [ENTER].

2. Choose “Set Peripheral Device Entry”, press [ENTER], then select “Event Trigger Operations” by pressing [ENTER]. The event trigger menu displays.

3. Select any of the monitoring elements by moving the cursor bar and pressing [ENTER] to enable or disable the association with related system events.

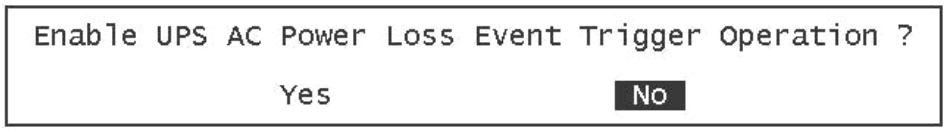

NOTE

The last condition, the “Temperature Threshold,“ is associated with a configurable time buffer before an automatic shutdown. Please refer to the next section for details.

11.2.2 Operation Theory

To reduce the chance of data loss due to hardware failure, the controller/subsystem automatically commences the following actions when a component failure is detected:

- Switches its caching mode from “write-back” to the conservative “write-through.”
- Flushes all cached data.
- Raises the rotation speed of cooling fans.

The Trigger

The mode-switching and cache-flush operations can be triggered by the occurrences of the following conditions:

- Controller failure (Dual-controller Models): a controller fails in a dual-redundant controller configuration, the surviving controller no longer has the protection of

synchronized cache by having the replica of unfinished writes in its partner.

- BBU low or failed: f a battery fails or is under-charge, the unfinished writes cannot be supported if power outage occurs.
- UPS AC power loss: Even with the buffer provided by the UPS, if power outage occurs, cached data should be immediately distributed to hard drives before the battery charge in UPS runs out.
- Power supply failure
- Fan failure
- Temperature exceeds threshold

If one or more of the event triggers listed above are enabled, the occurrence of the above conditions forces the controller/subsystem to adopt the “write-through” caching mode. Once the faulty condition is corrected, the controller/subsystem automatically restores the previous caching mode.

NOTE

The temperature thresholds refer to those set for both sensors on the RAID controller boards and those placed within the subsystem enclosure. In terms of the controller temperature, board 1 refers to the main circuit board and board 2 refers to the second-level I/O board or the daughter card. If any of the threshold values set for any sensor is exceeded, the reaction mode is automatically triggered.

If a battery is not installed in your RAID subsystem, the “BBU Low or Failed“ option should be disabled.

11.2.3 Auto Shutdown on Elevated Temperature

Go to: View and Edit Peripheral Devices > Set Peripheral Device Entry > Event Trigger Operations > Temperature Exceeds Threshold

System components can be damaged if operated under elevated temperature. You can configure the time periods between the detection of exceeded thresholds and

the controller's commencing an automatic shutdown.

The shutdown does not electrically disconnect the subsystem. When shutdown is commenced, the subsystem stops responding to I/O requests and flushes all cached writes in its memory. During that time, system administrators should have been notified of the condition and have begun restoring proper cooling of the subsystem. Extended operation under critical conditions like elevated temperature greatly reduces system efficiency and will eventually cause component failure.

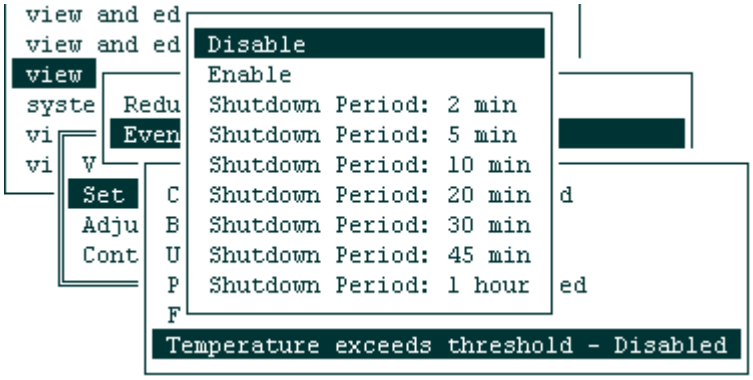

11.2.4 Voltage and Temperature Self-monitoring

Go to: View and Edit Peripheral Devices > Controller Peripheral Device Configuration > View Peripheral Device Status

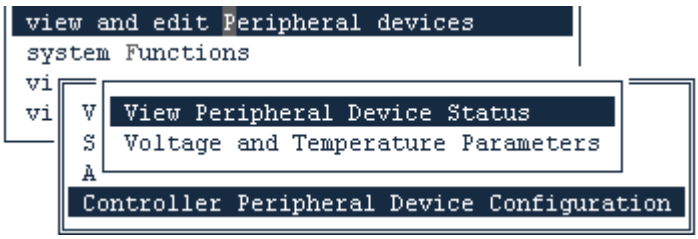

ITEM	VALUE	STATUS
+3.3V Value	3.368 V	Operation Normal
+5V Value	5.18 V	Operation Normal
+12V Value	12.32 V	Operation Normal
CPU Temperature	55.0 (C)	Within Safe Range
Controller Temperature(1)	42.5 (C)	Within Safe Range
Controller Temperature(2)	51.0 (C)	Within Safe Range

11.2.5 Changing Monitoring Thresholds

Go to: View and Edit Peripheral Devices > Controller Peripheral Device Configuration > Voltage and Temperature Parameters

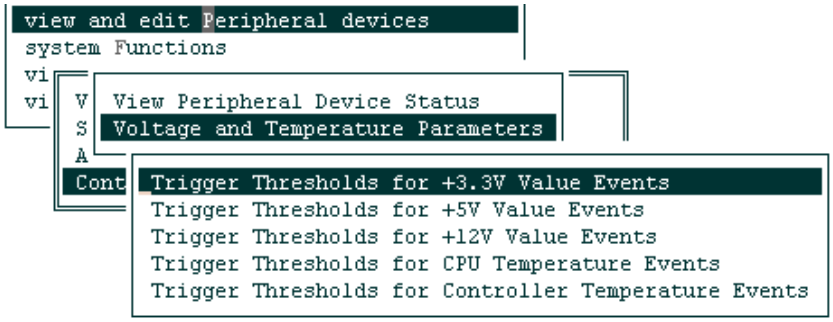
view and edit Peripheral devices
system Functions
View Peripheral Device Status
Voltage and Temperature Parameters
Trigger Thresholds for +3.3V Value Events
Trigger Thresholds for +5V Value Events
Trigger Thresholds for +12V Value Events
Trigger Thresholds for CPU Temperature Events
Trigger Thresholds for Controller Temperature Events

NOTE Do not change the threshold values unless you need to coordinate the RAID controller's values with that of your enclosure. If a value exceeding the safety range is entered, an error message will prompt and the new parameter will be ignored. For example, if the controller operates in a system enclosure where the upper limit on ambient temperature is relatively higher or lower, adjusting the default thresholds can coordinate the controller status monitoring with that of enclosure specifications.

1. Scroll down and select an item to configure.

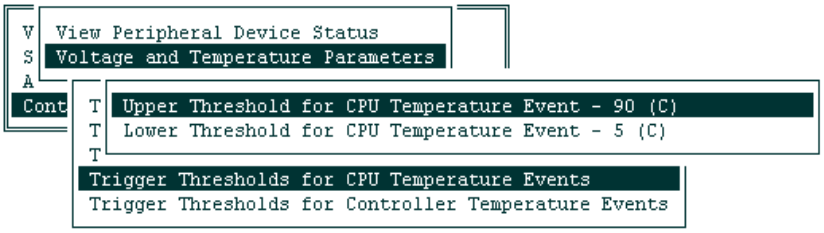
View Peripheral Device Status
Voltage and Temperature Parameters
Upper Threshold for CPU Temperature Event - 90 (C)
Lower Threshold for CPU Temperature Event - 5 (C)
Trigger Thresholds for CPU Temperature Events
Trigger Thresholds for Controller Temperature Events

2. Select an item, such as "Trigger Thresholds for CPU Temperature Events." Press [ENTER] and a list of selections will appear. You can change the upper or lower threshold values by keying a number. Press [ENTER] to confirm.

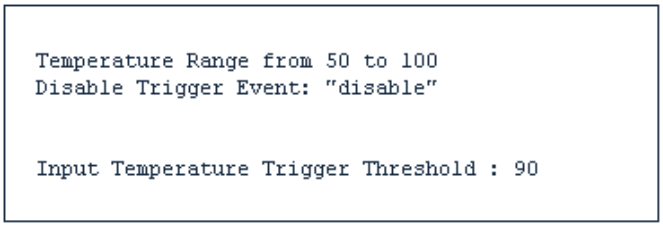

3. A configuration window will prompt. Enter any value within the safety range. Values exceeding the safety range will be rejected by controller firmware.

4. Follow the same method to modify other threshold parameters.

12 Data Integrity

This chapter discusses various firmware mechanisms that help to ensure data integrity.

No system is completely safe from hardware faults. For example, although the chance of occurrence is considerably low, the occurrences of bad blocks on two (RAID 5) or three (RAID 6) hard drives can fail a whole data set. When properly configured, the functions below help to minimize the chance of data loss:

NOTE

Some of the configuration options may not be available to all sub-revisions of firmware.

All figures in this chapter are showing examples of a management console over an RS-232 or telnet connection.

12.1 Restoring an Accidentally Deleted LD

Go to: View and Edit Logical Drives

If users accidentally delete a logical drive, the result is catastrophic. A Restore option is added to salvage an accidentally deleted LD. As long as the original member drives are not removed or configured into other logical drives, you can restore a deleted logical drive and bring it online.

If any of the original members is missing (not including a previously-failed member), you will not be able to restore a logical drive.

The members of a deleted LD will be indicated as “FRMT (formatted) drives” with array information still intact in its 256MB reserved space. These drives will not be converted into auto-hot-spares unless users manually put them into other uses.

Restoration Procedure:

1. Shown below is an empty index of a deleted LD (LG0) in the View and Edit Logical Drives menu.

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	O	C	#LN	#SB	#FL	NAME
0													
A1	14A48EC9	NA	RAID5	1430	GOOD	I		7	B	3	0	0	
A2	7A028143	NA	RAID5	1430	GOOD	I		7	B	3	0	0	
3													
4													
5													
6													
7													

2. Move cursor bar to an accidentally deleted LD, press the Space key on it.

3. The deleted LD will be shown. Press Enter on it.

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	O	C	#LN	#SB	#FL	NAME
A0	31178AE4	NA	RAID3	1430	GOOD			4	B	0	0	0	
A2	7A028143	NA	RAID5	1430	GOOD			7	B	3	0	0	
3													
4													
5													
6													
7													

4. The Restore command will prompt. Move cursor bar to it and press Enter.

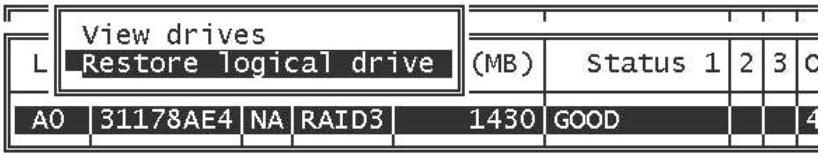

5. When prompted by the confirm box, choose Yes.

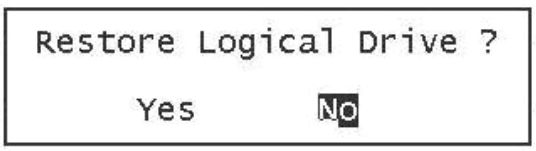

6. When restored, the recovered logical drive should appear in the LD list, and an event message showing the change of state.

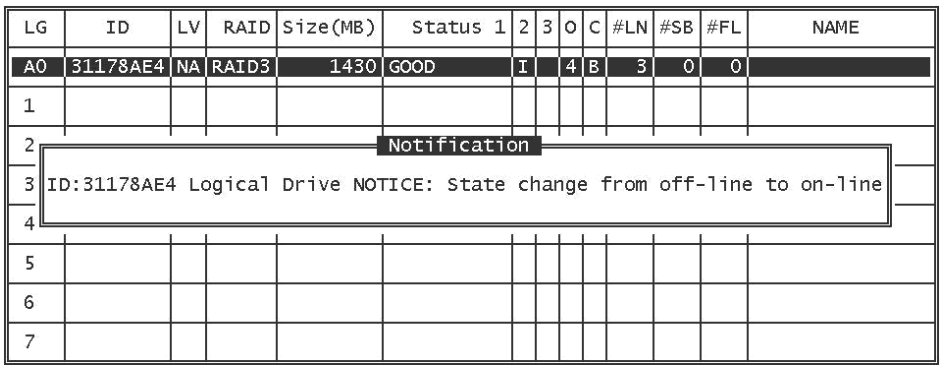
LG ID LV RAID Size(MB) Status 1 2 3 O C #LN #SB #FL NAME
A0 31178AE4 NA RAID3 1430 GOOD I 4 B 3 0 0
1
2
3
4
5
6
7
Notification
ID:31178AE4 Logical Drive NOTICE: State change from off-line to on-line

12.2 Detecting Failed Drives

12.2.1 Auto Rebuild on Drive Swap Check Time

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Auto Rebuild on Drive Swap

NOTE

When this option is enabled, it consumes a little portion of system resource by constantly checking drive buses.

The "Auto Rebuild on Drive Swap" timeout is enabled by choosing a time value. The RAID controller will poll all connected drives through the controller's drive channels at the assigned interval. Drive removal will be detected even if a host does not attempt to access data on that specific drive.

If the "Auto Rebuild on Drive Swap" timeout is set to "Disabled" (the default setting is "Disabled"), the controller will not be able to detect any drive removal that occurs after the controller initialization process. The controller will only be able to detect drive removal when host access is directed to the drive side.

The "Auto Rebuild on Drive Swap" check time is the interval at which the controller checks to see if a failed drive has been swapped. When a member of a logical drive fails, the controller will continuously scan the drive bus (at the selected time interval). Once the failed drive has been swapped with a drive that has the adequate capacity to rebuild the logical drive, the rebuild will begin automatically.

The default setting is "15 seconds," meaning that the controller will automatically scan the drive busses if a failed drive has been replaced. To change the timeout, select a time interval.

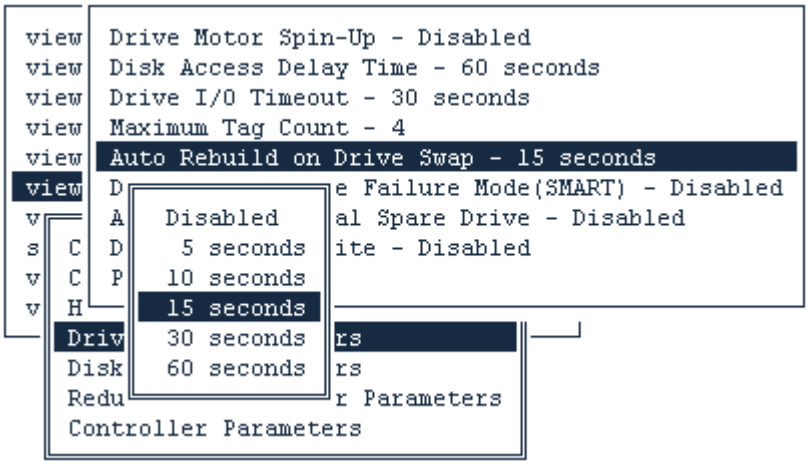

12.2.2 Auto-Assign Global Spare Drive

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Auto-Assign Global Spare Drive

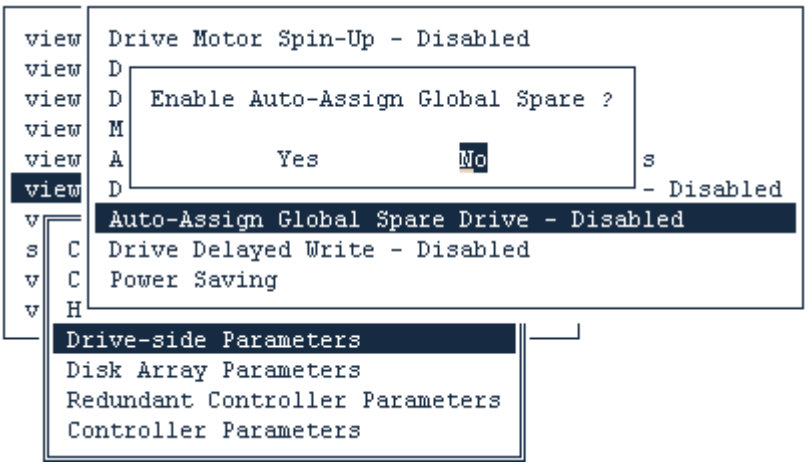

The “Auto-Assign” function automatically assigns any “new” drives that are not included in logical configurations as Global Spares.

NOTE

The Auto-Assign Global Spare applies to drive interfaces that support “auto detect,” such as Fibre Channel, SATA, and SAS interfaces. Disk drives of these interfaces can be detected shortly after they are installed.

Enabling the Function:

If a drive has a capacity smaller or apparently larger than the members of configured arrays, the controller may avoid using it as a global spare.

Enable the function and reset the controller for the configuration to take effect.

12.3 Scheduled Maintenance

12.3.1 Task Scheduler

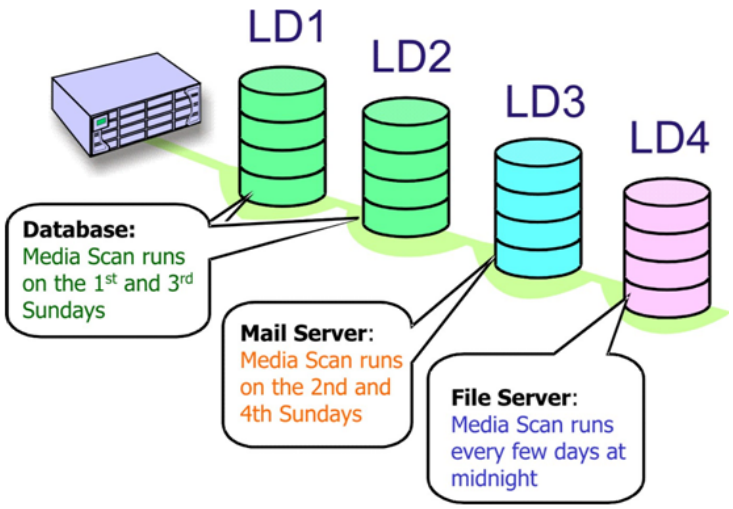

The Task Scheduler functionality allows Media Scans to be scheduled beginning at a specified start time and repeating at regular intervals defined by a configurable interval period. Each such schedule can be defined to operate on all drives of a certain class, all member drives of a specified logical drive, spare drives, or all member drives of all logical drives. UIs supported are the text-based utility accessed through RS-232C serial connection/telnet and RAIDWatch GUI manager.

The Task Scheduler allows firmware to automatically perform media scans on specific RAID arrays saving you the efforts to manually initiate the processes. Scans take place at a preferred time when the subsystem is less stressed by daily service, e.g., Sundays or midnight.

12.3.2 Setting Task Scheduler

Go to: View and Edit Logical Drives > Media Scan > Task Schedule

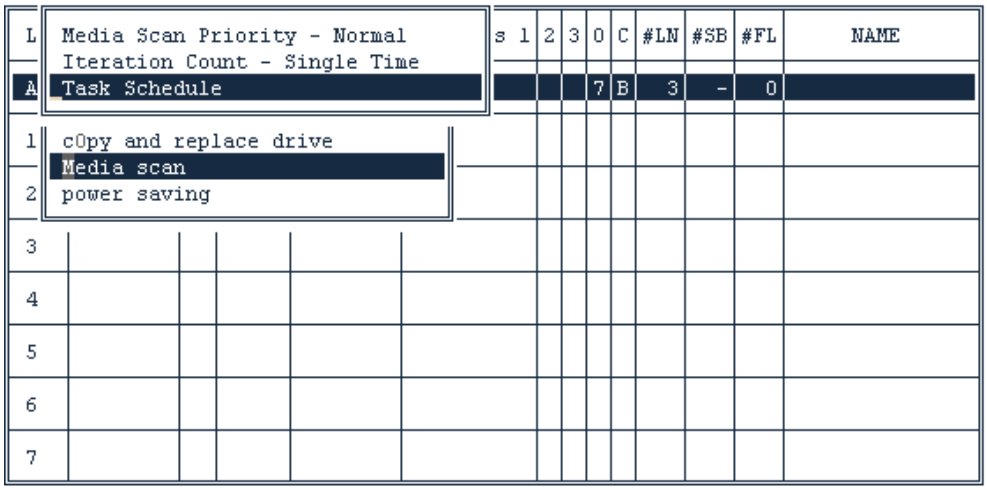

1. Select “Task Scheduler” by pressing [ENTER].

2. If there is no preset schedule, a confirm box will prompt.

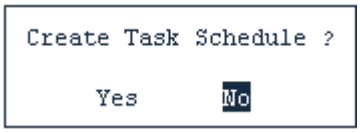

3. Press [ENTER] on an existing schedule to display the configuration options. You may choose to check information of a task schedule, to create a new schedule, or to remove a configured schedule.

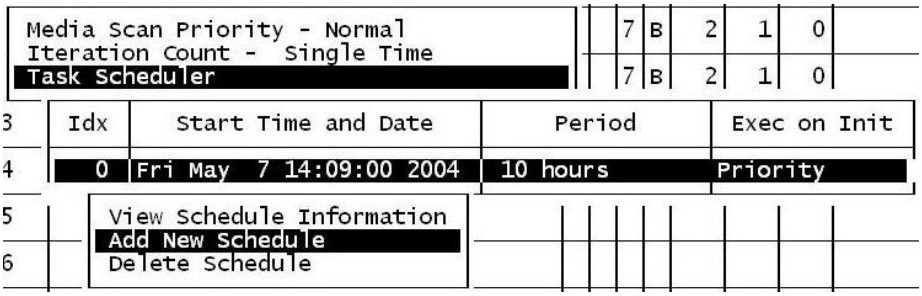

4. To configure a task schedule, browse through the following options and make necessary changes:

5. Execute on Controller Initialization: This option determines whether Media Scan is automatically conducted whenever the RAID system is reset or powered on.

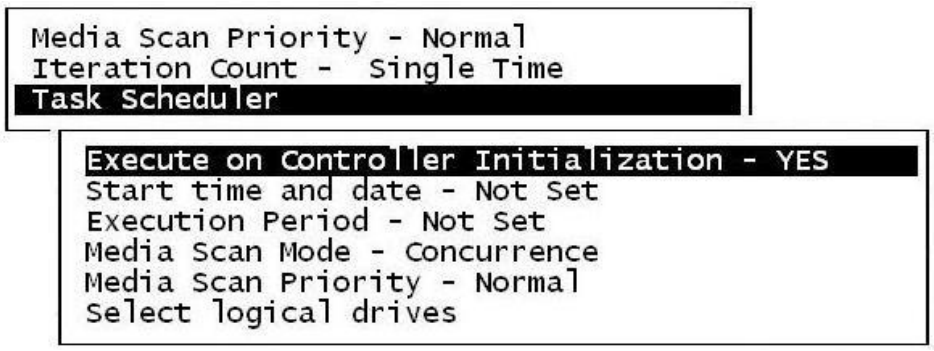
Media Scan Priority - Normal
Iteration Count - Single Time
Task Scheduler
Execute on Controller Initialization - YES
Start time and date - Not Set
Execution Period - Not Set
Media Scan Mode - Concurrence
Media Scan Priority - Normal
Select logical drives

6. Start Time and Date: Enter time and date in its numeric representatives in the following order: month, day, hour, minute, and the year.

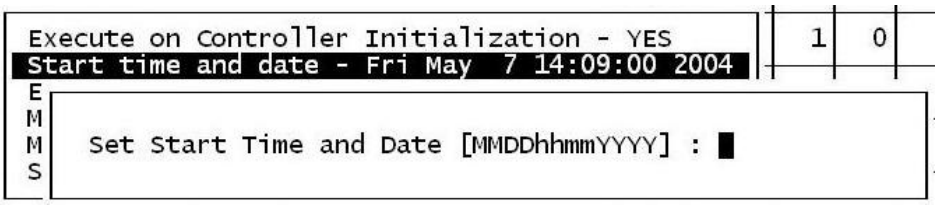
Execute on Controller Initialization - YES
Start time and date - Fri May 7 14:09:00 2004
1
0
Set Start Time and Date [MMDDhhmmYYYY] :

7. Execution Period: The scheduler memorizes the date and the time the actions are to be executed. Select one of the following:
   If the action is intended to be executed for one time only, select "Execution Once."
   In the case of a periodic action, the action is executed at the specified "start time," and then re-enacted at the time interval indicated in the execution period so as to be executed again later. The selectable interval ranges from one second to several weeks.

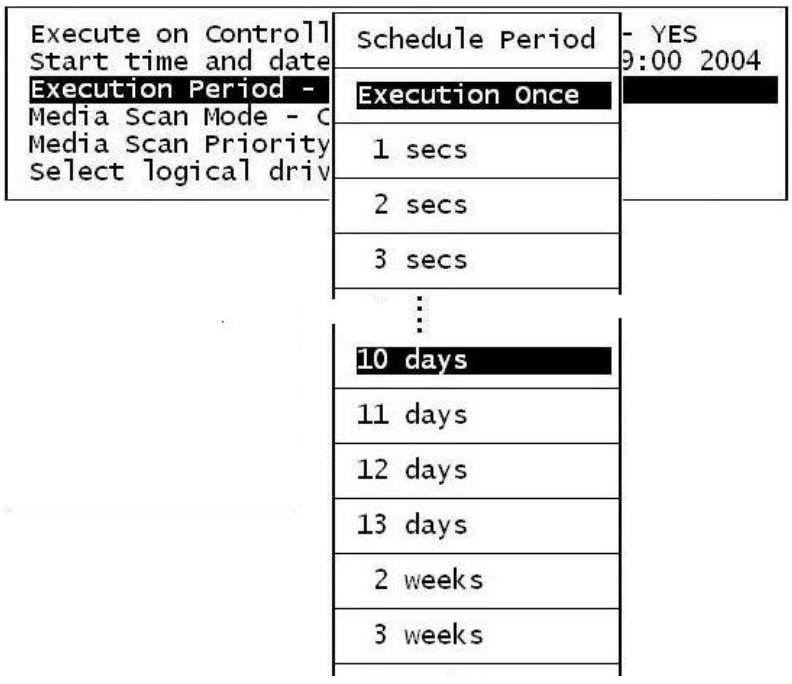

8. Media Scan Mode: If the maintenance schedule includes more than one logical drive, the scan can be performed simultaneously on multiple logical drives together or separately on one logical drive at a time following a sequential order.

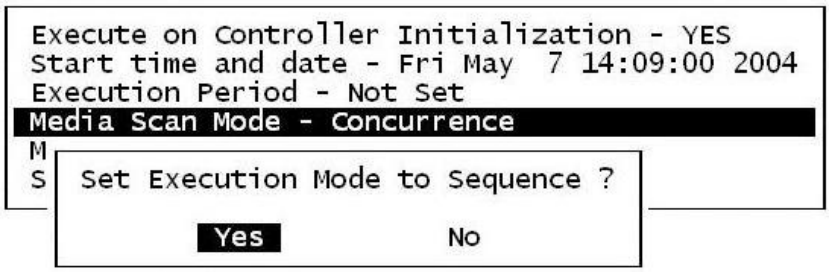
Execute on Controller Initialization - YES
Start time and date - Fri May 7 14:09:00 2004
Execution Period - Not Set
Media Scan Mode - Concurrence
Set Execution Mode to Sequence ?
Yes
No

9. Media Scan Priority: The scan priority determines how much of the system's resources will be consumed to perform the scheduled task. Select "Low" for better array performance and longer time to complete the media scan. Higher priority allows higher scan performance at the cost of reduced array performance.

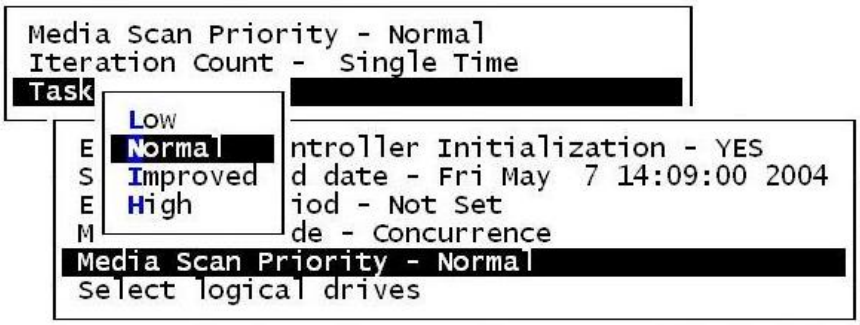

10. Select Logical Drives: Press [ENTER] on "Select Logical Drives" to bring out a sub-menu. From there you may include all configured arrays or press [ENTER] on "To Select Logical Drives" to select one or more specific logical drive(s).

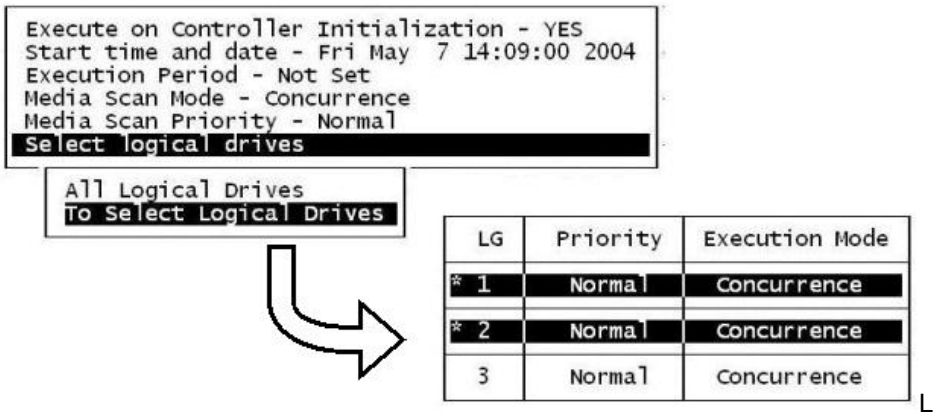

Logical drives can be tagged for inclusion by positioning the cursor bar on the logical drive and then pressing [ENTER]. An asterisk () mark will appear on the selected physical drive(s). To deselect the drive, press [ENTER] again on the selected drive. The "" mark will disappear. Use the same method to select more

logical drives if preferred.

When selection is done, press [ESC] to continue.

11. Confirming the Creation of a Task Schedule

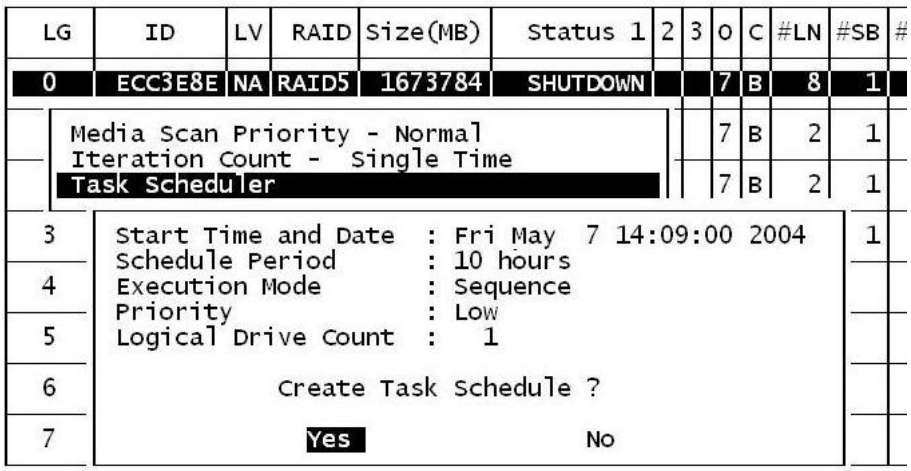
LG ID LV RAID Size(MB) Status 1 2 3 0 C #LN #SB #
0 ECC3E8E NA RAID5 1673784 SHUTDOWN 7 B 8 1
Media Scan Priority - Normal
Iteration Count - Single Time
Task Scheduler
7 B 2 1
7 B 2 1
Start Time and Date : Fri May 7 14:09:00 2004
Schedule Period : 10 hours
Execution Mode : Sequence
Priority : Low
Logical Drive Count : 1
Create Task Schedule ?
Yes
No

12. When finished with setting the scheduler options, press [ESC] to display a confirm box.

13. Verify all information in the box before choosing “Yes” to confirm and to complete the configuration process.

12.4 Manual Rebuild

If you want the controller to auto-detect a replacement drive, make sure you have a check time value set for the following option:

Auto Rebuild on Drive Swap check time

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Auto-Rebuild on Drive Swap

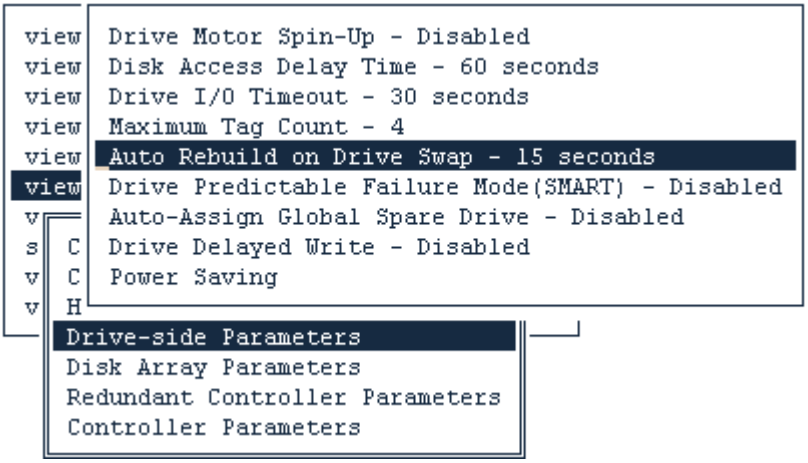
view Drive Motor Spin-Up - Disabled
view Disk Access Delay Time - 60 seconds
view Drive I/O Timeout - 30 seconds
view Maximum Tag Count - 4
view Auto Rebuild on Drive Swap - 15 seconds
view Drive Predictable Failure Mode(SMART) - Disabled
Auto-Assign Global Spare Drive - Disabled
Drive Delayed Write - Disabled
Power Saving
Drive-side Parameters
Disk Array Parameters
Redundant Controller Parameters
Controller Parameters

NOTE

A manual rebuild occurs in a system that has no hot-spare.

In a system configured with hot-spares, a rebuild should take place automatically.

The rebuild function will only appear if a logical drive (in RAID level 1, 3, 5, or 6) has a failed member.

Carefully verify the location of a failed drive before replacement takes place. Removing the wrong drive will fatally fail a logical drive and the data loss will occur.

1. Before physically replacing a failed drive, you should verify the messages as shown below:

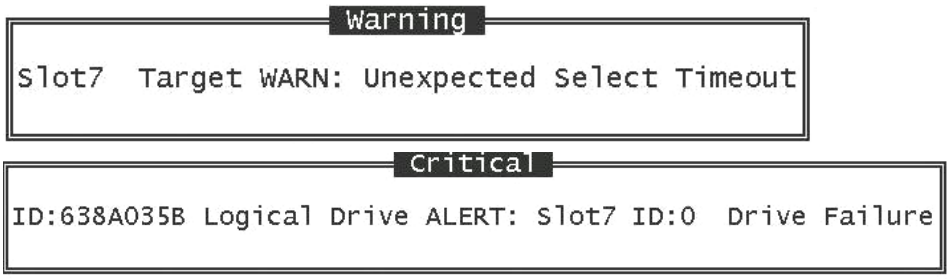
Warning
Slot7 Target WARN: Unexpected Select Timeout
Critical
ID:638A035B Logical Drive ALERT: Slot7 ID:0 Drive Failure

2. You should also check the logical drive member list in “View and Edit Logical Drives” -> “View drives.” The failed drive’s status should be indicated as “BAD.”

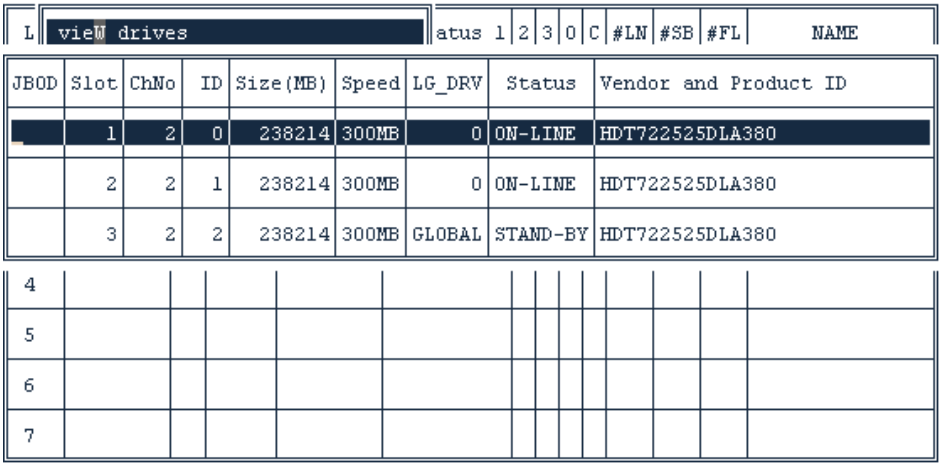
L
view drives
atus 1 2 3 0 C #LN #SB #FL NAME
JBOD Slot ChNo ID Size(MB) Speed LG_DRV Status Vendor and Product ID
1 2 0 238214 300MB 0 ON-LINE HDT722525DLA380
2 2 1 238214 300MB 0 ON-LINE HDT722525DLA380
3 2 2 238214 300MB GLOBAL STAND-BY HDT722525DLA380
4
5
6
7

3. Make sure you correctly identify the location of the failed drive and replace it with a new drive.

4. Return to the “View and Edit Logical Drives” menu. Press [ENTER] on it and you should find the “Rebuild” option.

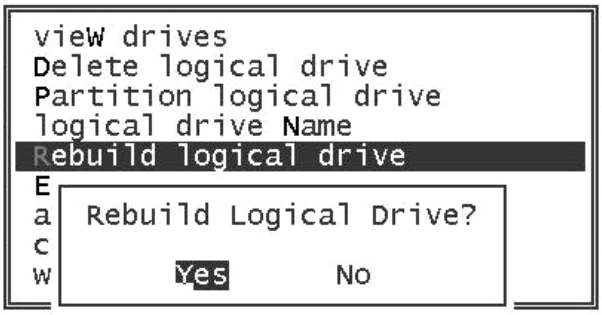

5. The rebuild should start. Press ESC to skip the message.

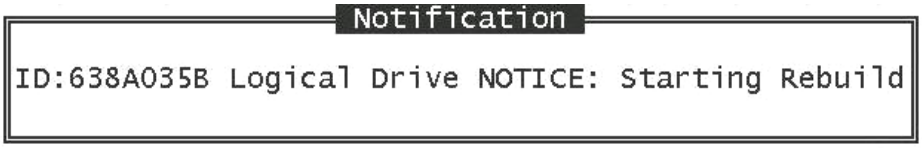

6. The rebuild progress will be indicated by a status bar.

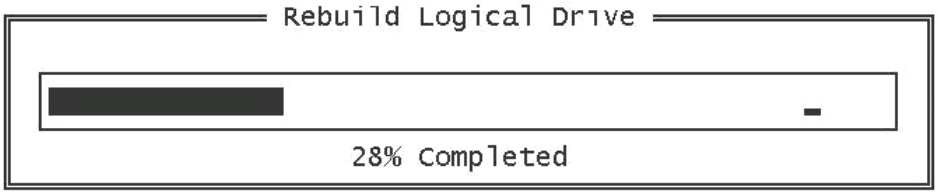

7. Upon the completion of rebuild, the following message will prompt. Press ESC to dismiss the message.

8. You may now return to the “View and Edit Logical Drives” menu to check if the array status is stated as “GOOD.”

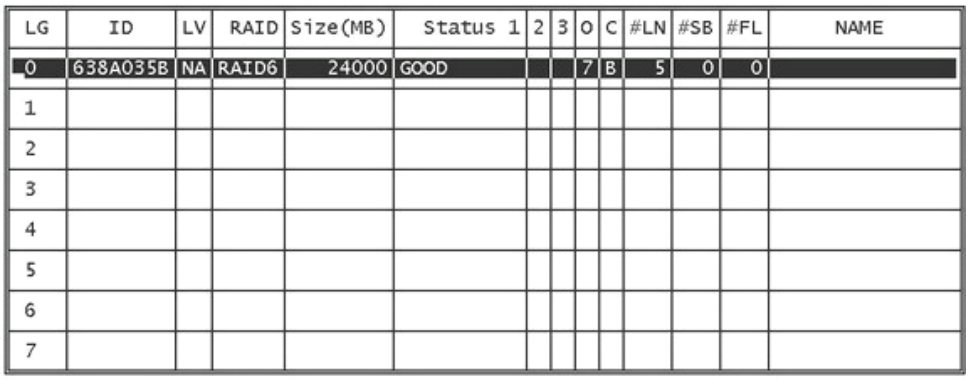

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	O	C	#LN	#SB	#FL	NAME
0	638A035B	NA	RAID6	24000	GOOD			7	B	5	0	0	
1													
2													
3													
4													
5													
6													
7													

12.5 Regenerating Logical Drive Parity

Go to: View and Edit Logical Drives > (Logical Drive) > Execute Regenerate Logical

Drive Parity

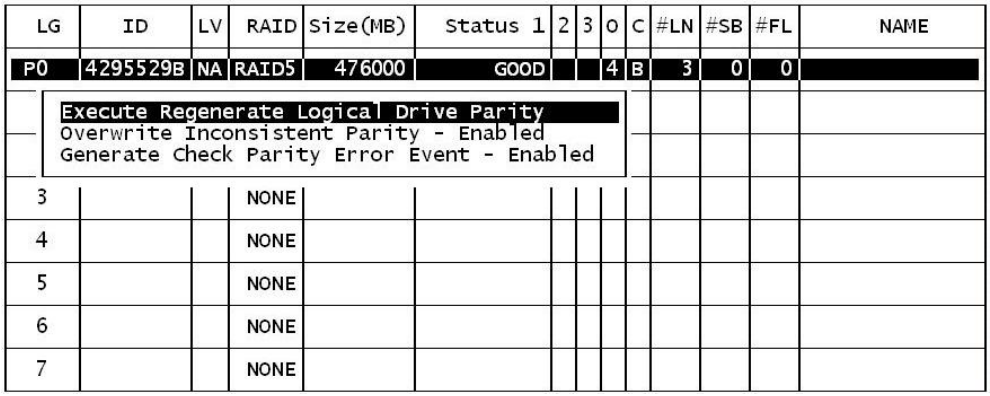
LG ID LV RAID Size(MB) Status 1 2 3 0 C #LN #SB #FL NAME
P0 4295529B NA RAID5 476000 GOOD 4 B 3 0 0
Execute Regenerate Logical Drive Parity
Overwrite Inconsistent Parity - Enabled
Generate Check Parity Error Event - Enabled
3 NONE
4 NONE
5 NONE
6 NONE
7 NONE

Parity regeneration is a function manually performed onto RAID-1/3/5/6 arrays to determine whether inconsistency has occurred with data parity.

12.5.1 Overwriting Inconsistent Parity

Go to: View and Edit Logical Drives > (Logical Drive) > Overwrite Inconsistent Parity

Default is “enabled.”

If an array’s data parity is seriously damaged, restoring parity data by regenerating and overwriting the original data may cause data loss. Disable this option if you suspect parity data has been seriously corrupted.

12.5.2 Generating Check Parity Error Event

Go to: View and Edit Logical Drives > (Logical Drive) > Generate Check Parity Error Event

Default is “enabled.”

When enabled, parity inconsistency will be reported as system events.

NOTE

If a regenerating process is stopped by a drive failure, the process cannot be restarted until the logical drive is successfully rebuilt by having its failed member replaced.

12.6 Setting Disk Array Parameters

12.6.1 Rebuild Priority

Go to: View and Edit Configuration Parameters > Disk Array Parameters

The system firmware provides a background rebuilding ability. This means firmware is able to serve I/O requests while rebuilding logical drives. The time required to rebuild a logical drive depends largely on the total capacity of the logical drive being rebuilt. Additionally, the rebuilding process is totally transparent to the host computer and its operating system. This option determines how much resources will be utilized when rebuilding a logical drive. For example, setting to Low allows system to have more elbow room to continue ongoing services while the rebuild will take a longer time to complete.

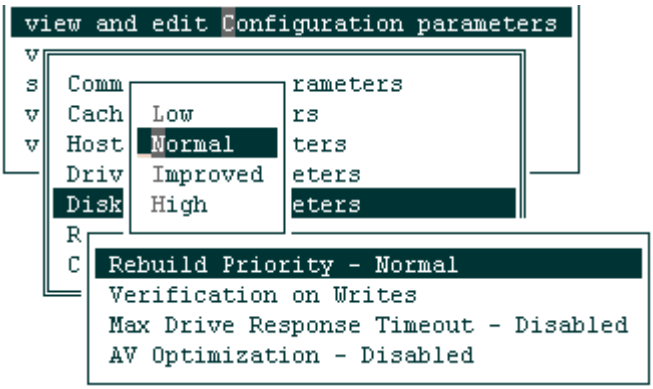

12.6.2 Verification on Writes

Go to: View and Edit Configuration Parameters > Disk Array Parameters > Verification on Writes

Errors may occur when a hard drive writes data. To avoid the write error, the controller can force hard drives to verify written data. There are three selectable methods:

- Verification on LD Initialization Writes: Performs Verify-after-Write when initializing a logical drive
- Verification on LD Rebuild Writes: Performs Verify-after-Write during the rebuild process
- Verification on LD Normal Drive Writes: Performs Verify-after-Write during normal I/Os

Each method can be enabled or disabled individually. Hard drives will perform Verify-after-Write according to the selected method.

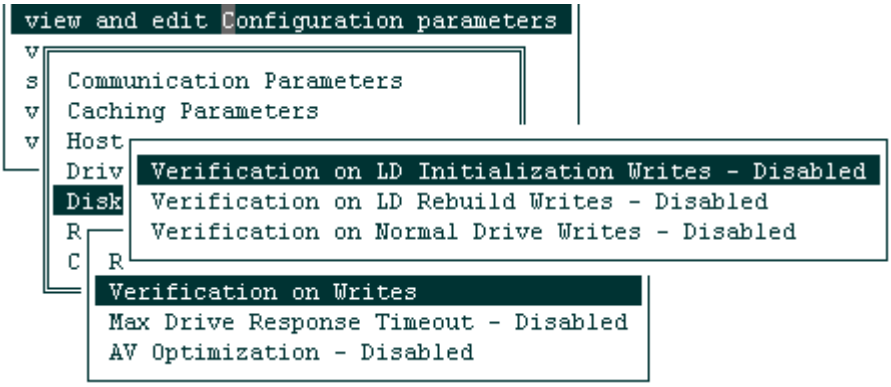

1. Move the cursor bar to the desired item, then press [ENTER].

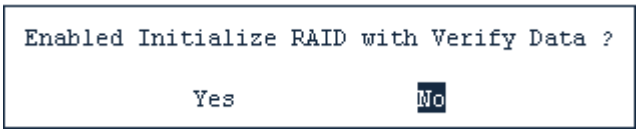

2. Choose Yes in the confirm box to enable or disable the function. Follow the same procedure to enable or disable each method.

NOTE

The “verification on Normal Drive Writes” method will bring overhead to the “write” performance of your RAID system.

13 Array Expansion

The array expansion functions allow you to expand storage capacity without the cost of buying new equipment. Expansion can be completed on-line while the system is serving host I/Os.

13.1 Overview

13.1.1 What is RAID Expansion and how does it work?

Before the invention of RAID Expansion, increasing the capacity of a RAID system meant backing up all data in the disk array, re-creating the disk array configuration with new drives, and then restoring data back into system.

Galaxy's RAID Expansion technology allows users to expand a logical drive by adding new drives, or replacing drive members with drives of larger capacity. Replacing is done by copying data from the original members onto larger drives; the smaller drives can then be replaced without powering down the system.

13.1.2 Notes on Expansion

Expansion Capacity:

When a new drive is added to an existing logical drive, the capacity brought by the new drive appears as a new partition. For example, if you have 4 physical drives (36GB each) in a logical drive, and each drive's maximum capacity is used, the capacity of the logical drive will be 108GB. (One drive's capacity is used for parity, e.g., RAID 3). When a new 36GB drive is added, the capacity will be increased to 144GB in two separate partitions (one is 108GB and the other 36GB).

Size of the New Drive:

A drive used for adding capacity should have the same or more capacity as other drives in the array.

Applicable Arrays:

Expansion can only be performed on RAID 0, 1, 3, 5, and 6 logical drives. Expansion cannot be performed on logical configurations that do not have parity, e.g., NRAID or RAID 1.

NOTE

Expansion on RAID0 is not recommended, because the RAID0 array has no redundancy. Interruptions during the expansion process may cause unrecoverable data loss.

Interruption to the Process:

Expansion should not be canceled or interrupted once begun. A manual restart should be conducted after the occurrence of a power failure or interruption of any kind.

13.1.3 Expand Logical Drive: Re-striping

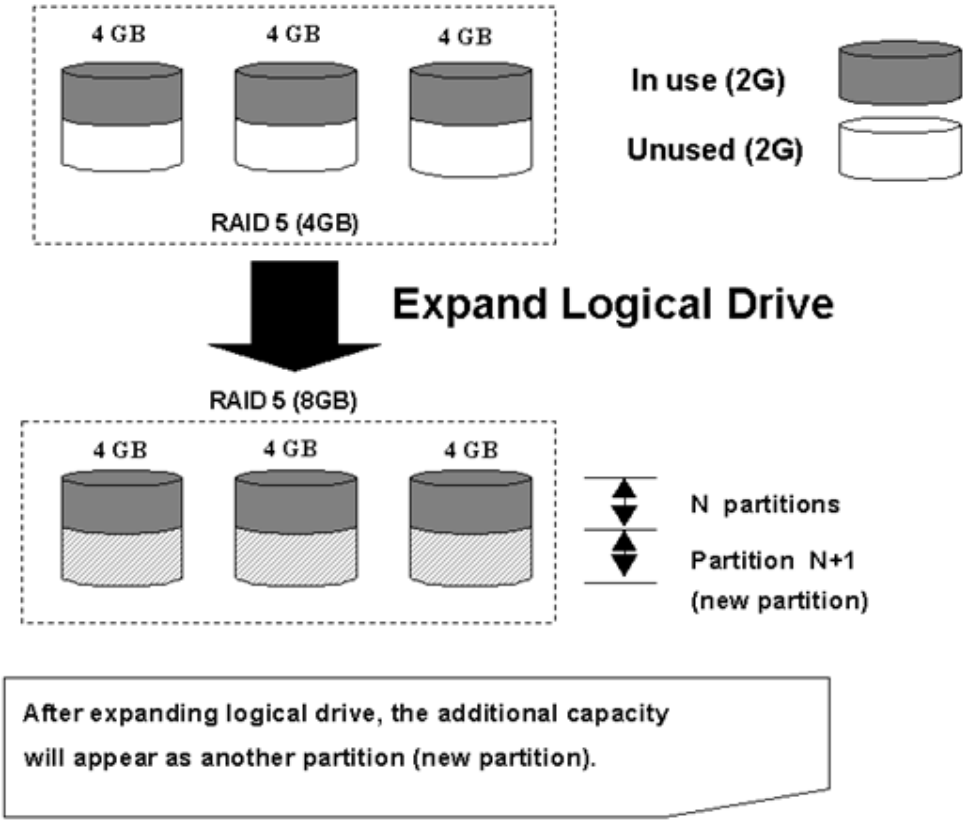

RAID levels supported: RAID 0, 1, 3, 5 and 6

Expansion can be performed on logical drives or logical volumes under the following conditions:

- There is unused capacity in a logical unit
- Capacity is increased by using member drives of larger capacity (see Copy and

Replace in the discussion below)

Data is recalculated and distributed to drive members or members of a logical volume. Upon the completion of the process, the added or the previously unused capacity will become a new partition. The new partition must be made available through host LUN mapping in order for a host adapter to recognize its presence.

13.2 Mode 1 Expansion: Adding Drives to a Logical Drive

13.2.1 Overview

Use drives with the same capacity as the original drive members. Once completed, the added capacity will appear as another partition (new partition). Data is automatically re-striped across the new and old members during the add-drive process. See the diagram below to get a clear idea:

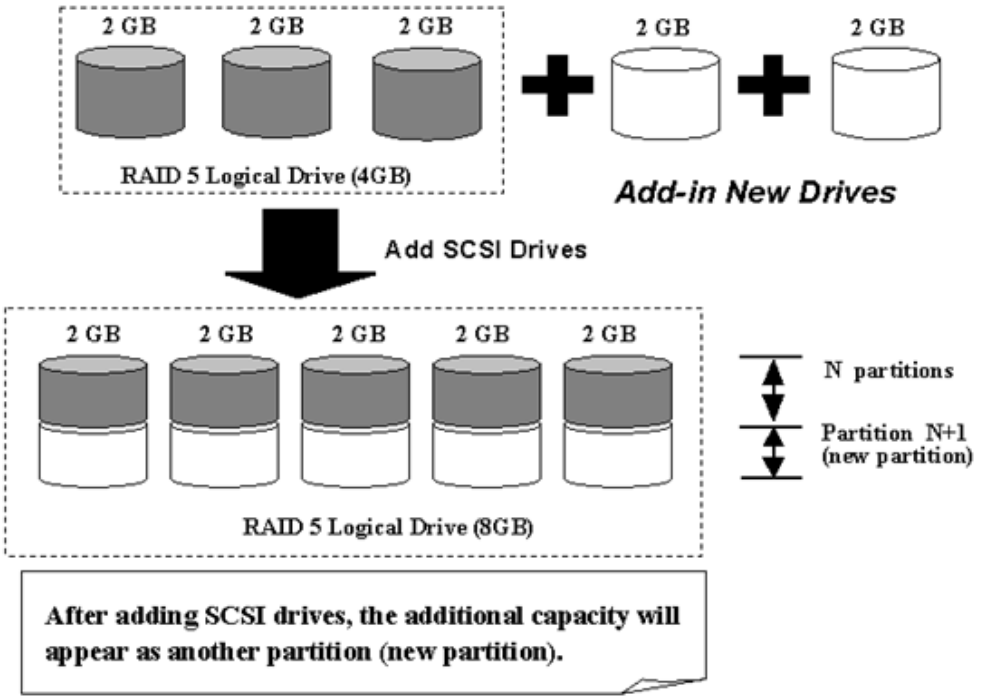

RAID levels supported: RAID 0, 1, 3, 5, and 6. The new partition must be made available through a host ID/LUN.

13.2.2 Add Drive Procedure

Go to: View and Edit Logical Drives > Add Drives

NOTE

The drive selected for adding should have a capacity no less than the original member drives. If possible, use drives of the same capacity because all drives in the array are treated as though they have the capacity of the smallest member in the array.

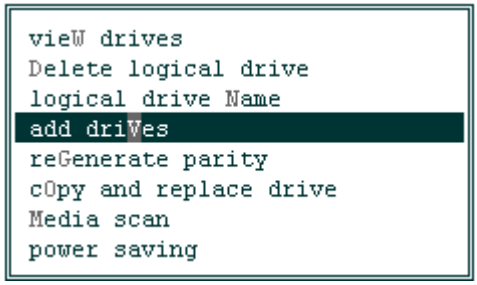

1. Available drives will be listed. Select one or more disk drive(s) to add to the target logical drive by pressing [ENTER]. The selected drive will be indicated by an asterisk "*" mark.

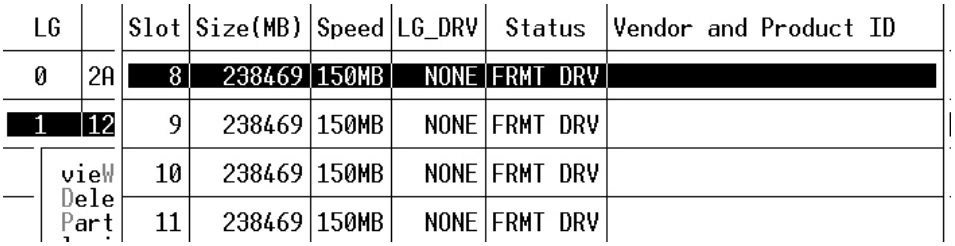

2. Press [ESC] to proceed and the notification will prompt.

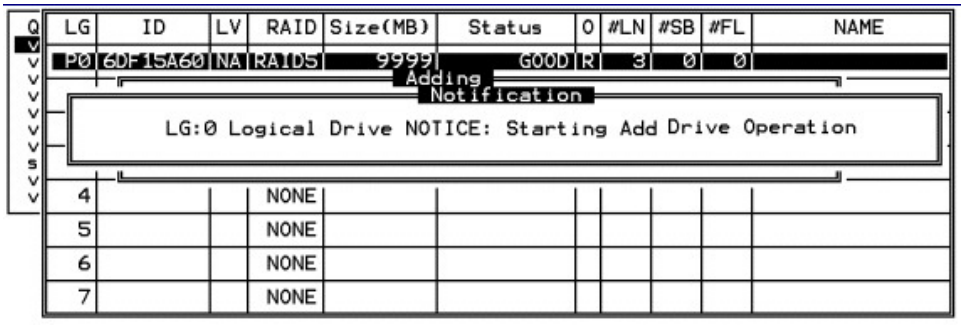

3. Press [ESC] again to cancel the notification prompt; a status bar will indicate the percentage of progress.

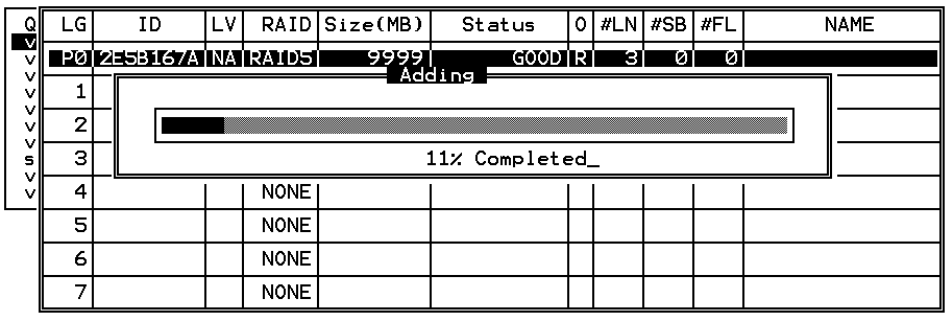

4. Upon completion, there will appear a confirming notification. The capacity of the added drive will appear as an unused partition.

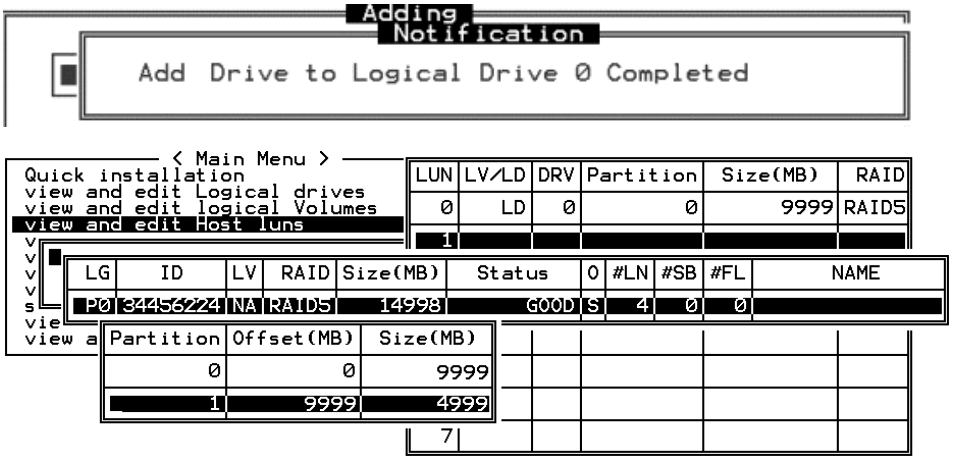
Adding
Notification
Add Drive to Logical Drive 0 Completed
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
LUN LV/LD DRV Partition Size(MB) RAID
0 LD 0 0 9999 RAID5
LG ID LV RAID Size(MB) Status 0 #LN #SB #FL NAME
P0 34456224 NA RAID5 14998 GOOD S 4 0 0
Partition Offset(MB) Size(MB)
0 0 9999
1 9999 4999

5. The added capacity will be automatically included, meaning that you do not have to "expand logical drive" later. Map the added capacity to another host ID/LUN to make use of it.

As diagrammed above, in "View and Edit Host LUN," the original capacity is 9999MB, its host LUN mapping remains unchanged, and the added capacity appears as the second partition.

NOTE

Expansion by adding drives can not be canceled once started. If power failure occurs, the expansion will be paused and the controller will NOT restart the expansion when power comes back on. Resumption of the RAID expansion must be performed manually.

If a member drive of the logical drive fails during RAID expansion, the expansion will be paused. The expansion will resume after the logical drive rebuild is completed.

13.3 Mode 2 Expansion: Copy and Replace Drives with Drives of Larger Capacity

13.3.1 Overview

You may also expand your logical drives by copying and replacing all member drives with drives of higher capacity. Please refer to the diagram below for a better understanding. The existing data in the array is copied onto the new drives, and then the original members can be removed.

When all the member drives have been replaced, execute the “Expand Logical Drives” function to make use of the expansion capacity.

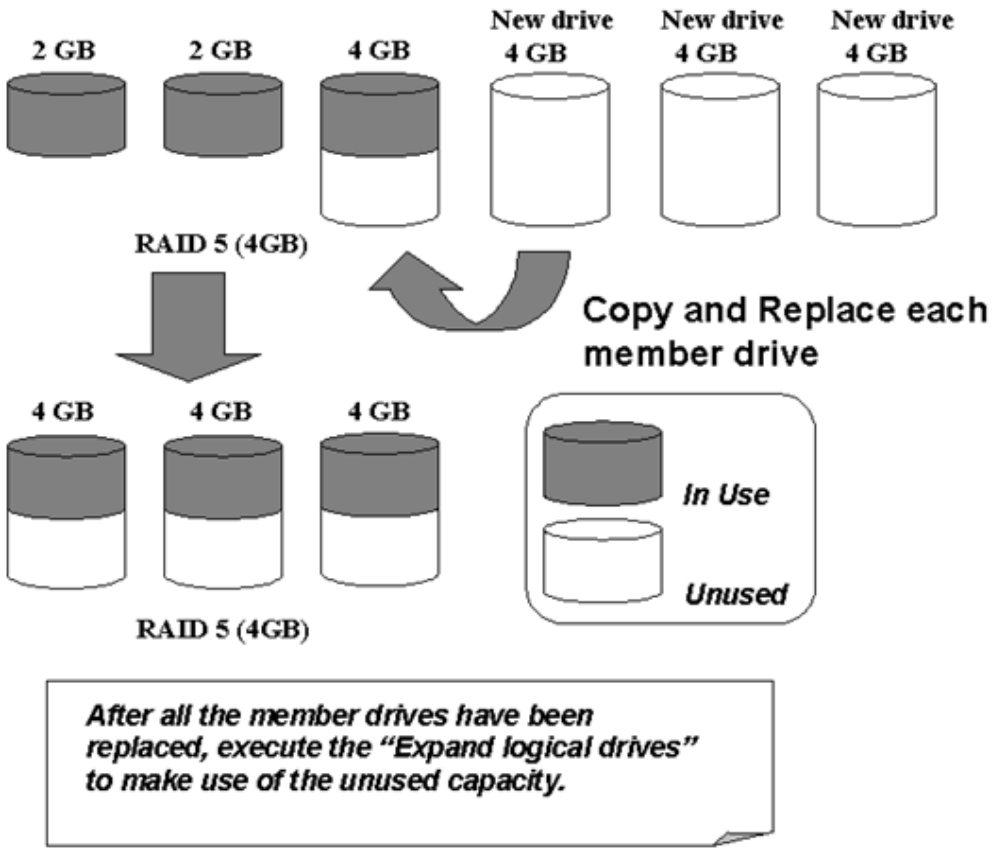

RAID levels supported: RAID 0, 1, 3, 5 and 6

13.3.2 Copy and Replace Procedure

Go to: View and Edit Logical Drives > (Logical Drive) > Copy and Replace Drive

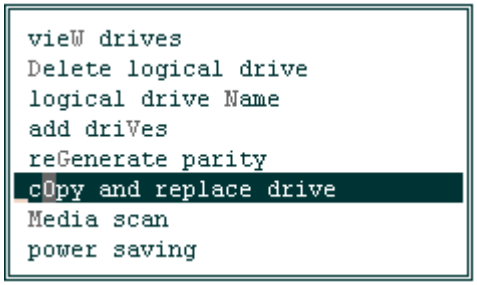

1. The array members will be listed. Select the member drive (the source drive) you want to replace with a larger one.

2. Select one of the members as the "source drive" (status indicated as ON-LINE) by pressing [ENTER]; a table of available drives will prompt.

3. Select a "new drive" to copy the capacity of the source drive to. The channel number and ID number of both the “Source Drive” and the “Destination Drive” will be indicated in the confirm box.

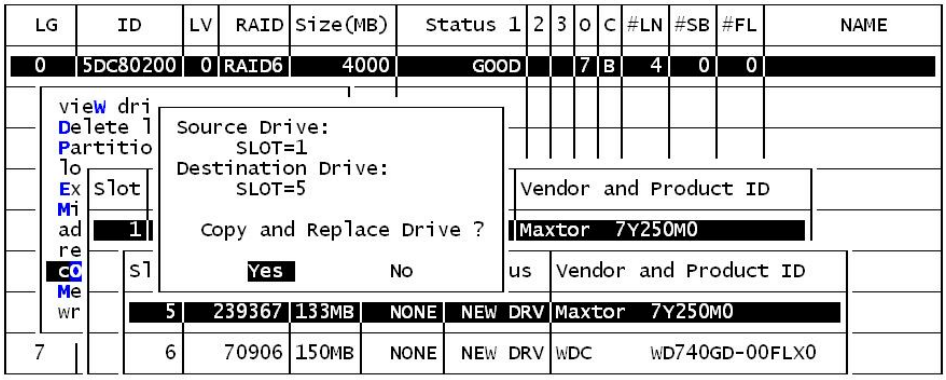
LG ID LV RAID Size(MB) Status 1 2 3 0 C #LN #SB #FL NAME
0 5DC80200 0 RAID6 4000 GOOD 7 B 4 0 0
Source Drive:
SLOT=1
Destination Drive:
SLOT=5
Copy and Replace Drive ?
Yes
No
Slot
Vendor and Product ID
1
Maxtor 7Y250M0
5 239367 133MB NONE NEW DRV Maxtor 7Y250M0
6 70906 150MB NONE NEW DRV WDC WD740GD-00FLX0

4. Choose Yes to confirm and proceed.

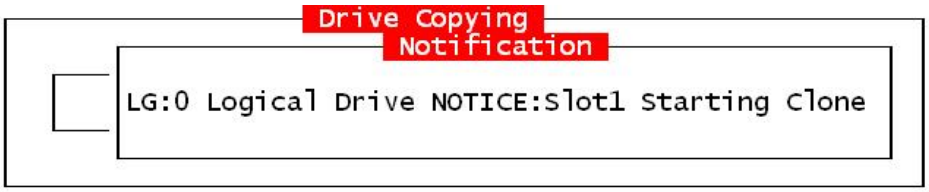

5. Press [ESC] to view the progress.

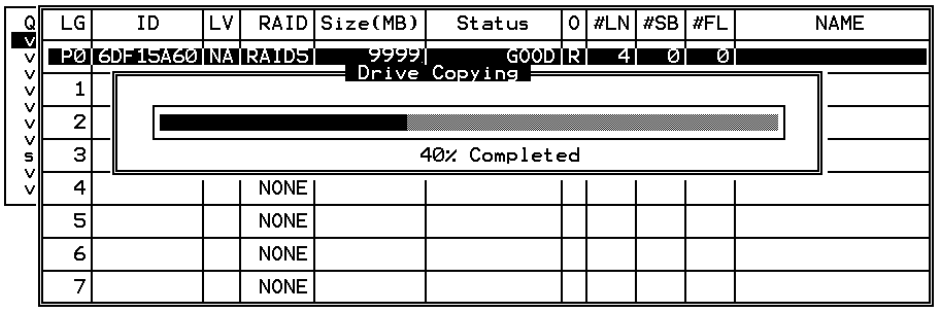

6. Completion of the Copy and Replace process will be indicated by a notification

message. Follow the same method to copy and replace every member drive. You may now perform “Expand Logical Drive” to make use of the added capacity, and then map the additional capacity to a host LUN.

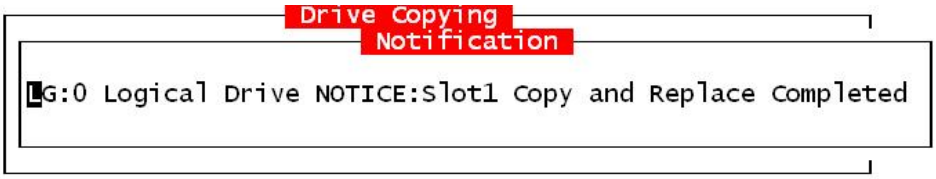

13.4 Expanding Logical Drives and Volumes

In the following example, the logical drive is originally composed of three member drives and each member drive has the capacity of 1GB. “Copy and Replace” has been performed on the logical drive and each of its member drives has been replaced by a new drive with the capacity of 2GB. The next step is to perform “Expand Logical Drive” to utilize the additional capacity brought by the new drives.

13.4.1 Expanding Logical Drives

Go to: View and Edit Logical Drives > (Logical Drive) > Expand Logicval Drive

1. Proceed by pressing [ENTER] or entering any value no larger than the "maximum drive expand capacity" and press [ENTER].

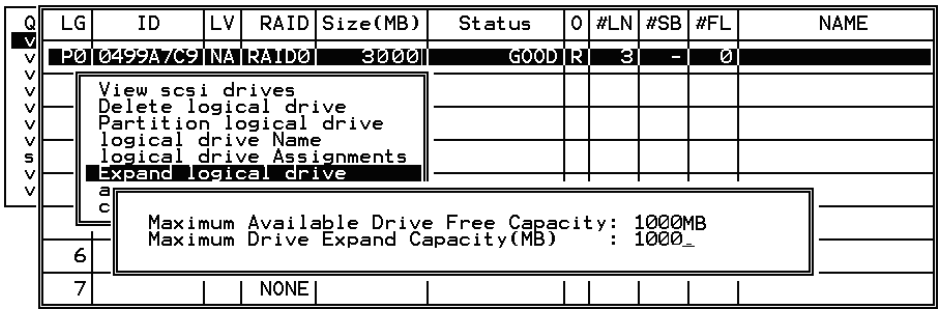
LG ID LV RAID Size(MB) Status O #LN #SB #FL NAME
P0 0499A7C9 NA RAID0 3000 GOOD R 3 - 0
View scsi drives
Delete logical drive
Partition logical drive
logical drive Name
logical drive Assignments
Expand logical drive
Maximum Available Drive Free Capacity: 1000MB
Maximum Drive Expand Capacity(MB) : 1000
NONE

2. Choose Yes to confirm.

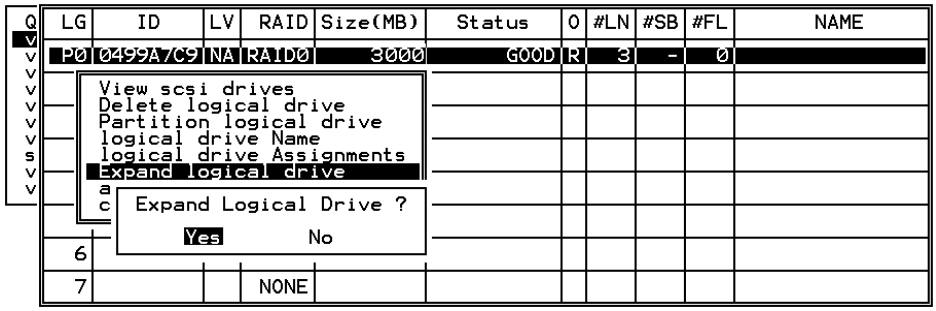
LG ID LV RAID Size(MB) Status O #LN #SB #FL NAME
P0 0499A7C9 NA RAID0 3000 GOOD R 3 - 0
View scsi drives
Delete logical drive
Partition logical drive
logical drive Name
logical drive Assignments
Expand logical drive
Expand Logical Drive ?
Yes No
NONE

3. Upon completion, you will be prompted by the notification message.

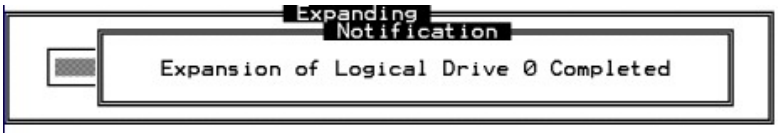

4. Press [ESC] to return to the previous menu screen.

5. As shown below, the total capacity of logical drive has been expanded to 6GB.

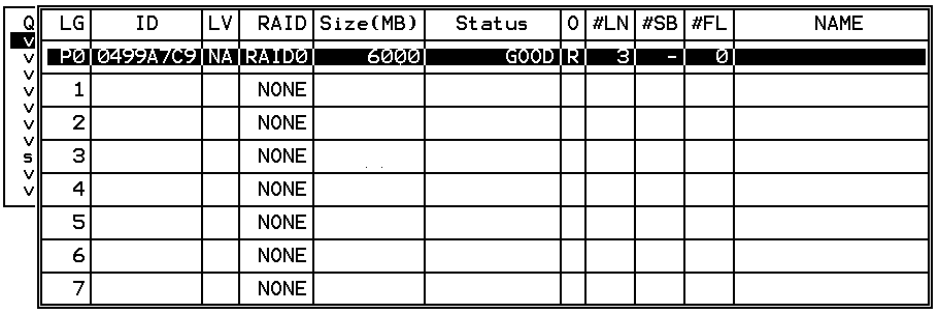

LG	ID	LV	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
P0	0499A7C9	NA	RAID0	6000	GOOD	R	3	-	0	
1			NONE							
2			NONE							
3			NONE							
4			NONE							
5			NONE							
6			NONE							
7			NONE							

13.4.2 Expanding Logical Volumes

> NOTE
>
> If the logical drive that has an expanded capacity is a member of a logical volume, make sure you expand all logical drives within the logical volume. A logical volume is made of logical drives that are ”striped” together. Unless all logical drives within a logical volume have excessive capacity, you cannot expand a logical volume.

1. To expand a logical volume, expand its logical drive member(s) and then perform “Expand logical volume.”

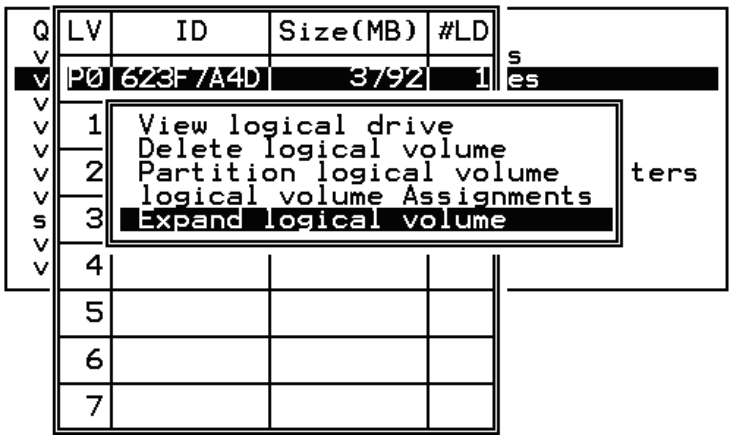

2. When prompted by "Expand Logical Volume?", choose Yes to confirm and the process will be completed immediately.

13.5 Configuration Example: Volume Extension in Windows®

The following demonstration will show how to expand the capacity of an existing logical drive by adding new drives and make it available to Windows file system.

Scenario

There is one logical drive made of two physical drives and it has been recognized by the Computer Management utility in Windows 2008 Server environment. A storage manager can increase volume capacity by adding additional drive(s) to it (if there are unused drives in the enclosure).

13.5.1 Prerequisites

In order to extend a logical volume, the following items need to be in place:

- At least one logical drive
- At least one logical volume including the abovementioned logical drive
- At least one partition inside the logical volume, mapped to the host

In Windows Server, open the Server Manager program and select Disk Management. You should be able to see the existing logical volume recognized along the physical drives (Disk 1 in below case).

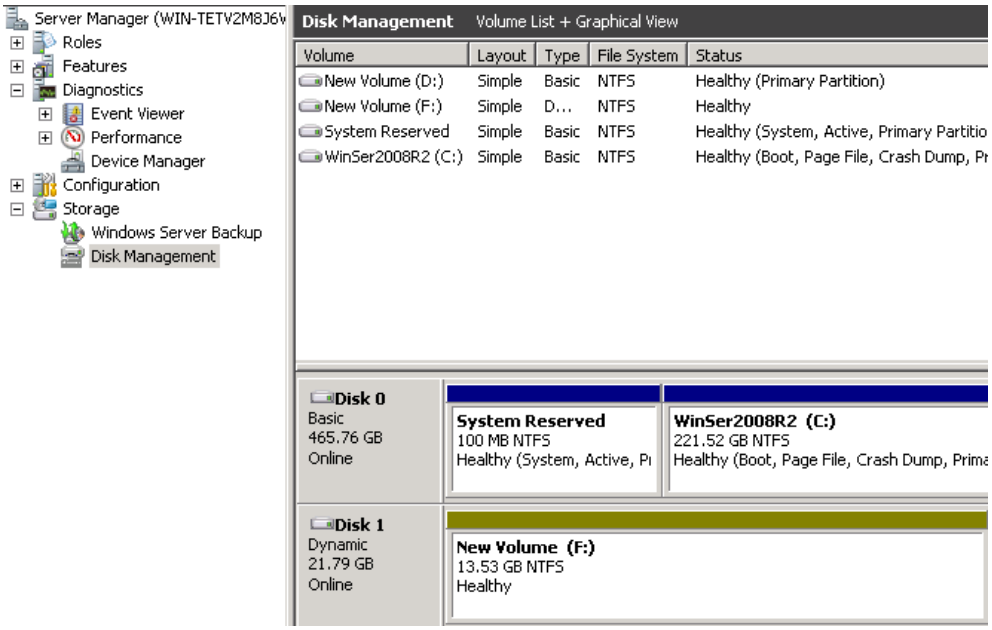

In the following procedures, we will follow these steps:

1. Expand the logical drive
2. Expand the logical volume
3. Expand the Partition
4. Reflect the expanded logical volume status in Windows.

13.5.2 Step 1: Expanding the Logical Drives

Go to: View and Edit Logical Drive > (Logical Drive) > Add Drives

In this example, we will expand the logical drive by adding a drive disk. You may also choose the Expand Logical Drive option.

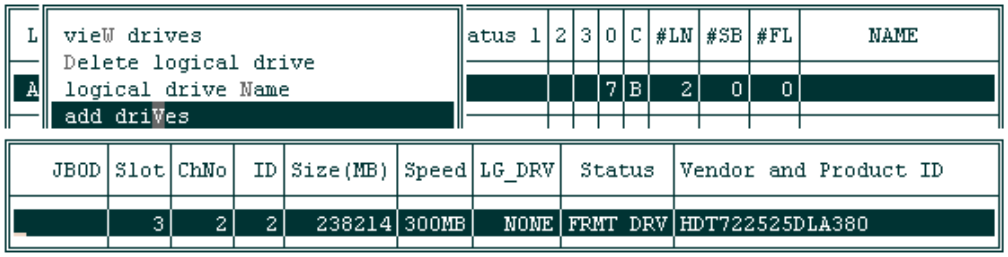
L
vieW drives
Delete logical drive
A
logical drive Name
add driVes
atus 1 2 3 0 C #LN #SB #FL NAME
7 B 2 0 0
JBOD Slot ChNo ID Size(MB) Speed LG_DRV Status Vendor and Product ID
3 2 2 238214 300MB NONE FRMT DRV HDT722525DLA380

1. Press Enter. The “JBOD” parameter changes to 1.

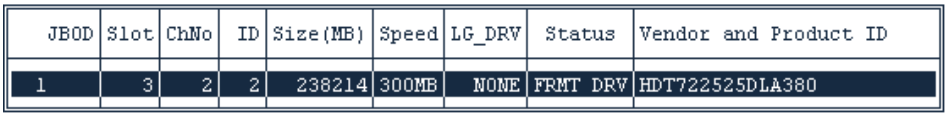

2. Press Esc. When the prompt appears, select Yes.

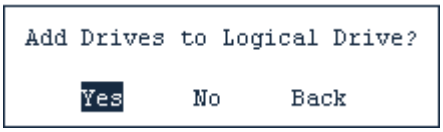

3. Wait until the “addition” is 100% done.

4. Repeat these steps for the rest of the logical drives included in the logical volume.

NOTE

You need to expand all logical drives in the target logical volume.

13.5.3 Step 2: Expanding the Logical Volume

Go to: View and Edit Logical Volumes > (Logical Volume) > Expand Logical Volume

NOTE

You cannot expand the logical volume until all the logical drive expansion process (including adding drive disks) have completed.

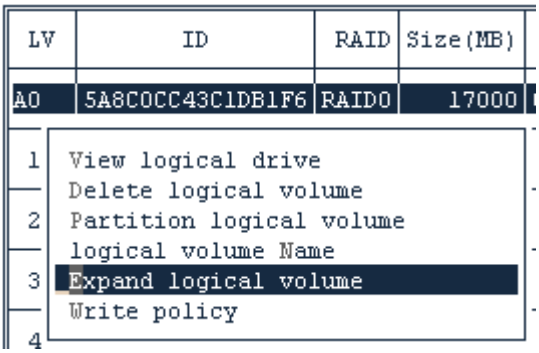

1. Specify the amount of expansion and press [ENTER].

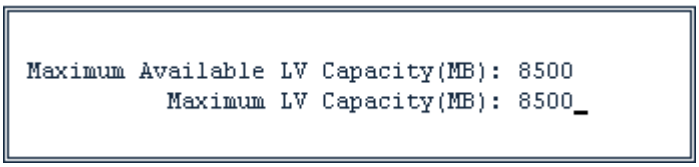

2. Select Yes to expand the logical volume.

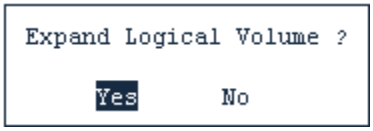

13.5.4 Step 3: Expanding the Partition

Go to: View and Edit Logical Volumes > (Logical Volume) > Partition Logical Volume > (Partition) > Expand Partition

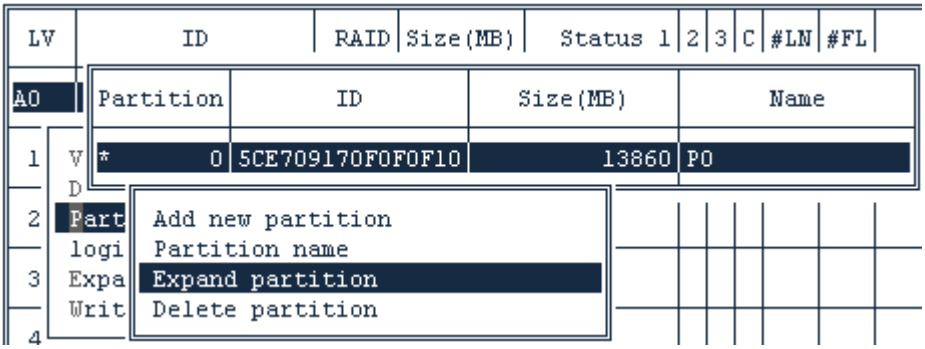

1. Select the capacity and press [ENTER].

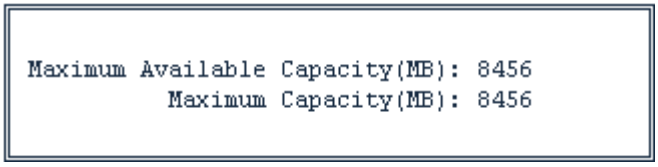

2. Select Yes.

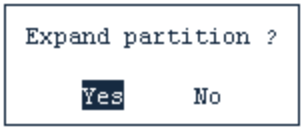

13.5.5 Step 4: Expand the Original Logical Volume in Computer Management Utility

1. Now return to the Window's Computer Management utility to view the additional disk.

2. Right-click on Disk Management and select ReScan.

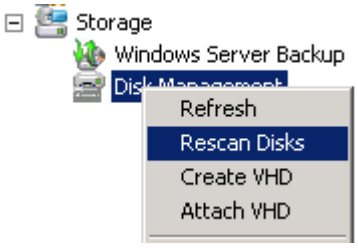

3. The newly added Logical Volume Partition will appear as a new unallocated disk space in Disk 1.

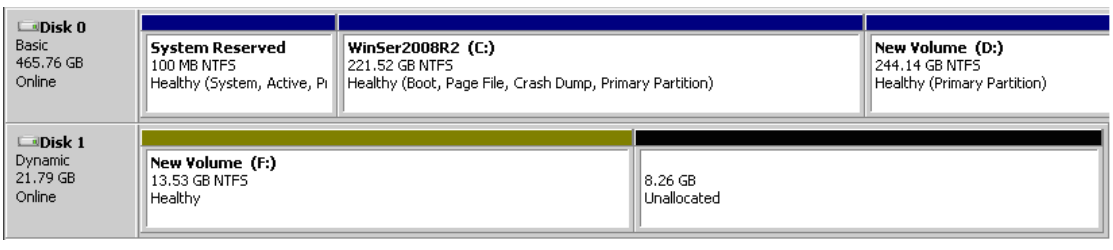

Open the Command Window and type in diskpart.exe. We will use Windows Diskpart utility to extend the existing partition without destroying the original data.

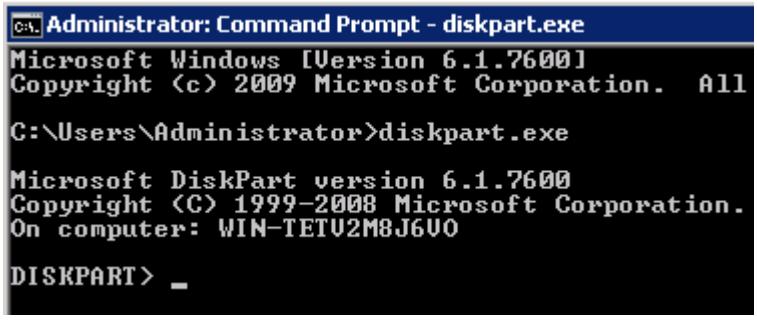

4. Type *list volume* to display the existing volumes on the computer.

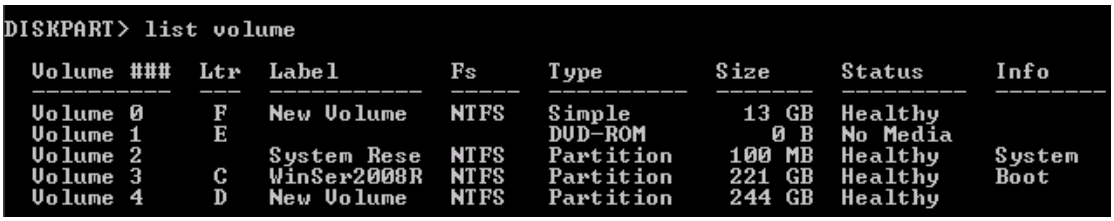
DISKPART> list volume
Volume ### Ltr Label Fs Type Size Status Info
Volume 0 F New Volume NTFS Simple 13 GB Healthy
Volume 1 E DVD-ROM 0 B No Media
Volume 2 System Rese NTFS Partition 100 MB Healthy System
Volume 3 C WinSer2008R NTFS Partition 221 GB Healthy Boot
Volume 4 D New Volume NTFS Partition 244 GB Healthy

5. Type *Select volume volume number* where volume number is number of the volume that you want to extend.

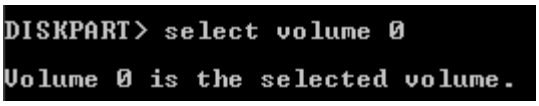

6. Type extend [size=n] [disk=n] [noerr]. The following describes the parameters:

   size=n The space, in megabytes (MB), to add to the current partition. If you do not specify a size, the disk is extended to use all the next contiguous unallocated space.

   disk=n The dynamic disk on which to extend the volume. Space equal to size=n is allocated on the disk. If no disk is specified, the volume is extended on the current disk.

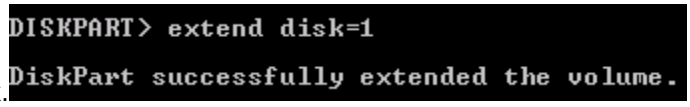

7. Now you can see the expanded volume.

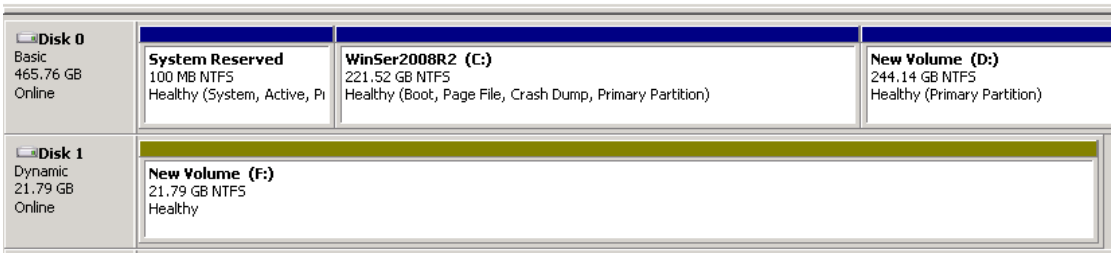

14 S.M.A.R.T. Support

With the maturity of technologies like S.M.A.R.T., drive failures can be predicted to certain degree. Before S.M.A.R.T., receiving notifications of drive bad block reassignments may be the most common omen that a drive is about to fail. In addition to the S.M.A.R.T.-related functions as will be discussed later, a system administrator can also choose to manually perform “Clone Failing Drive” on a drive which is about to fail.

This function provides system administrators a choice on when and how to preserve data from a failing drive. Although not necessary under normal conditions, you may also replace any drive at-will even when the source drive is healthy.

The “Clone Failing Drive” can be performed under the following conditions:

- Replacing a failing drive either detected by S.M.A.R.T. or notified by the controller.
- Manually replacing and cloning any drive with a new drive.

14.1 Cloning Failing Drive

Unlike the similar functions combined with S.M.A.R.T., the “Clone Failing Drive” is a manual function. There are two options for cloning a failing drive: “Replace after Clone” and “Perpetual Clone.”

14.1.1 Replacing After Clone

Go to: View and Edit Drives > (Drive) > Clone Failing Drive > Replace After Clone

Data on the source drive, the drive with predicted errors (or any selected member drive), will be cloned to a standby spare and replaced later by the spare. The status of the replaced drive, the original member drive with predicted errors, will be redefined as a “used drive.” System administrators may replace the “used drive” with a new one, and then configure the new drive as a spare drive.

1. Select “Replace After Clone.” The controller will automatically start the cloning process using the existing “stand-by” (dedicated/global spare drive) to clone the source drive (the target member drive with predicted errors). If there is no standby drive (local/global spare drive), you need to add a new drive and

configure it as a standby drive.

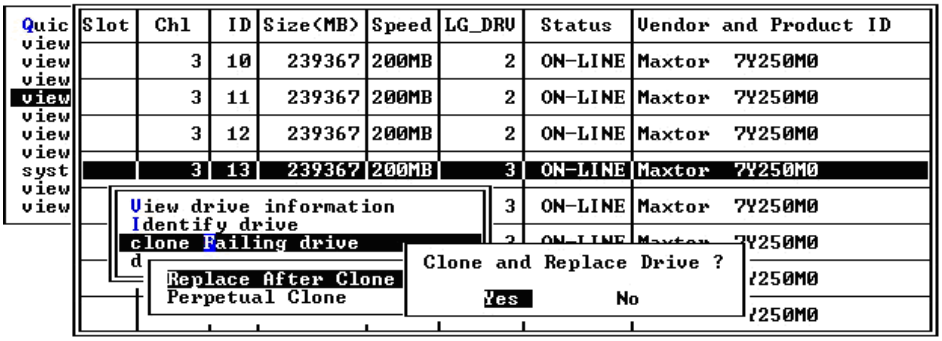
Quic Slot Chl ID Size(MB) Speed LG_DRV Status Vendor and Product ID
view 3 10 239367 200MB 2 ON-LINE Maxtor 7Y250M0
view 3 11 239367 200MB 2 ON-LINE Maxtor 7Y250M0
view 3 12 239367 200MB 2 ON-LINE Maxtor 7Y250M0
syst 3 13 239367 200MB 3 ON-LINE Maxtor 7Y250M0
view 3 ON-LINE Maxtor 7Y250M0
View drive information
Identify drive
clone Failing drive
Replace After Clone
Perpetual Clone
Clone and Replace Drive ?
Yes
No

2. The cloning process will begin with a notification message. Press [ESC] to proceed.

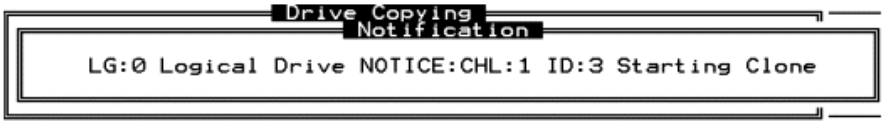

3. The cloning process will be indicated by a status bar.

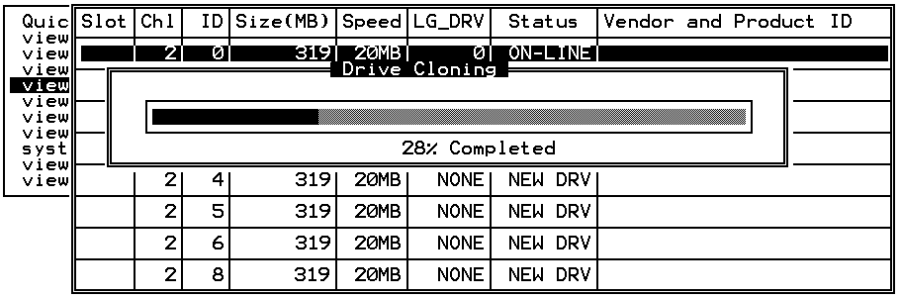
Quic Slot Chl ID Size(MB) Speed LG_DRV Status Vendor and Product ID
2 0 319 20MB 0 ON-LINE
Drive Cloning
28% Completed
2 4 319 20MB NONE NEW DRV
2 5 319 20MB NONE NEW DRV
2 6 319 20MB NONE NEW DRV
2 8 319 20MB NONE NEW DRV

4. You may quit the status bar by pressing [ESC] to return to the table of the connected drives. Select the drive indicated as “CLONING” by pressing [ENTER].

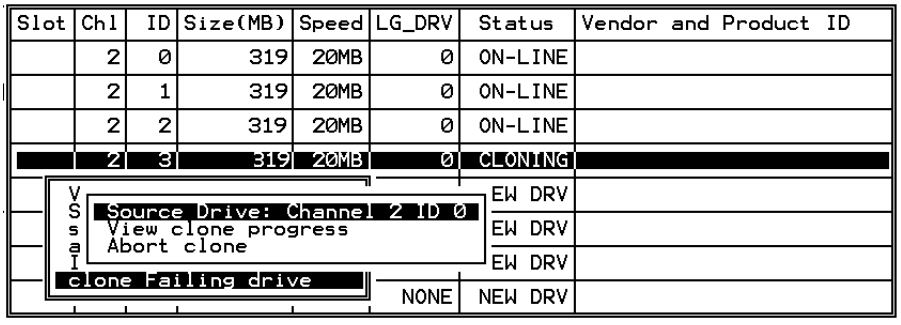
Slot Chl ID Size(MB) Speed LG_DRV Status Vendor and Product ID
2 0 319 20MB 0 ON-LINE
2 1 319 20MB 0 ON-LINE
2 2 319 20MB 0 ON-LINE
2 3 319 20MB 0 CLONING
Source Drive: Channel 2 ID 0
View clone progress
Abort clone
clone Failing drive
NONE NEW DRV

5. Select “Clone Failing Drive” again to view the current status. You may identify the source drive and choose to “View Clone Progress,” or “Abort Clone” if you happen to have selected the wrong drive.

6. When the process is completed, you will be notified by the following message.

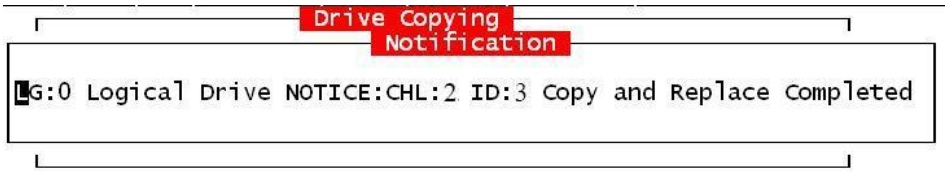

14.1.2 Perpetual Clone

Go to: View and Edit Drives > (Drive) > Clone Failing Drive > Perpetual Clone

The standby spare will clone the source drive, the member drive with predicted errors or any selected drive, without substituting it. The status of the spare drive will be displayed as “clone drive” after the cloning process. The source drive will remain a member of the logical drive. If the source drive fails, the clone drive can readily take its place in the array.

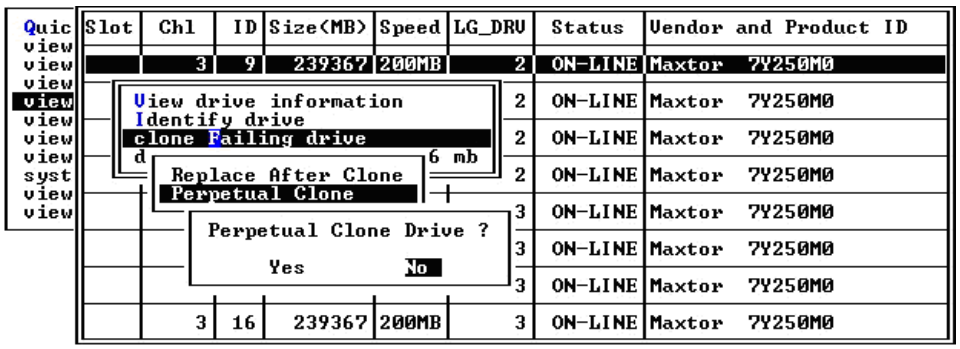
Slot Chl ID Size(MB) Speed LG_DRV Status Vendor and Product ID
3 9 239367 200MB 2 ON-LINE Maxtor 7Y250M0
View drive information
Identify drive
clone Failing drive
Replace After Clone
Perpetual Clone
Perpetual Clone Drive ?
Yes No
3 16 239367 200MB 3 ON-LINE Maxtor 7Y250M0

1. The controller will automatically start the cloning process using the existing “stand-by” (local/global spare drive) to clone the source drive (the target member drive).

2. The cloning process will begin with a notification message:

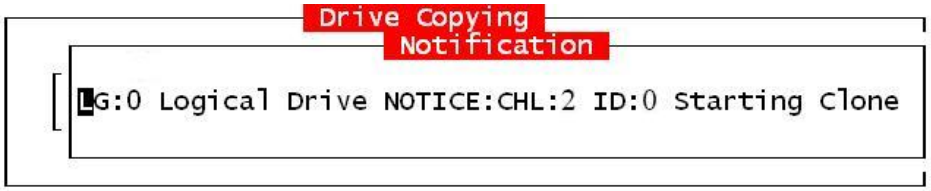

3. Press [ESC] to view the current progress:

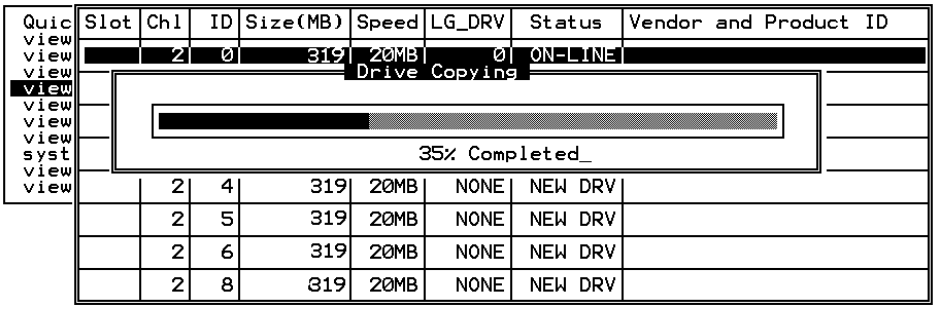

4. You may quit viewing the status bar by pressing [ESC] to return to the previous menu. Select the drive indicated as “CLONING” by pressing [ENTER]. Select “Clone Failing Drive” again to view the progress. You may identify the source drive and choose to “View Clone Progress” or “Abort Clone” if you happen to have selected the wrong drive.

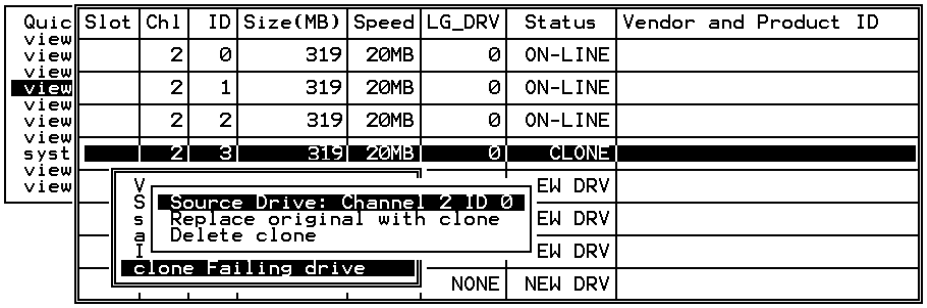

5. The cloning progress will be completed by a notification message as displayed below:

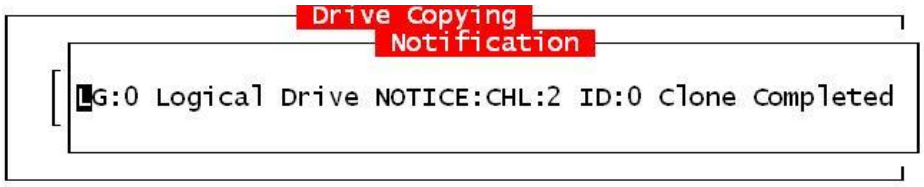

6. You may press [ESC] to clear the notification message to see the drives’ status after the cloning process. The source drive (Channel 1 ID 5) remains as a member of logical drive “0,” and the “stand-by” drive (Channel 1 ID 2, the dedicated/global spare drive) has become a “CLONE” drive.

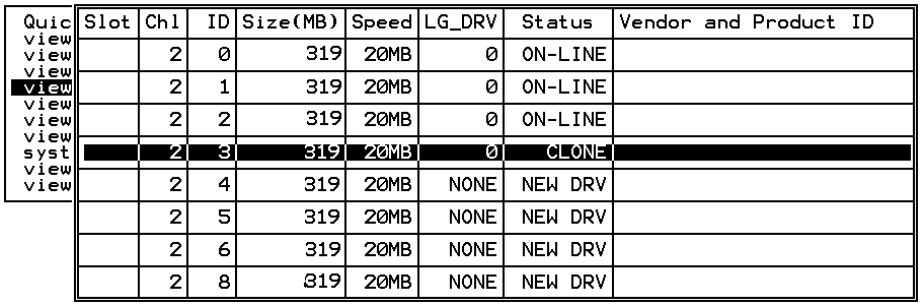

14.2 S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology)

This section provides a brief introduction to S.M.A.R.T. as one way to predict drive failure and Galaxy's implementations with S.M.A.R.T. for preventing data loss caused by drive failure.

14.2.1 Introduction

Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.) is an emerging technology that provides near-term failure prediction for disk drives. When S.M.A.R.T. is enabled, the drive monitors predetermined disk drives attributes that are susceptible to degradation over time.

If a failure is likely to occur, S.M.A.R.T. makes a status report available so that the host can prompt the user to backup data from the failing drive. However, not all failures can be predicted. S.M.A.R.T. predictions are limited to the attributes the drive can monitor which are selected by the device manufacturer based on the attribute's ability to contribute to predict degrading or fault conditions.

Although attributes are drive specific, a variety of typical characteristics can be identified:

- Head flying height
- Data throughput performance
- Spin-up time
- Re-allocated sector count
- Seek error rate
- Seek time performance
- Spin try recount
- Drive calibration retry count

Drives with reliability prediction capability only indicate whether the drive is "good" or "failing." In a SCSI environment, the failure decision occurs on the disk drive and the host notifies the user for action. The SCSI specification provides a sense bit to be flagged if the disk drive determines that a reliability issue exists. The system

then alerts the user/system administrator.

14.2.2 Galaxy's Implementations with S.M.A.R.T.

Galaxy uses the ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190 standard.

There are four selections related to the S.M.A.R.T. functions in firmware:

Disabled

Disables S.M.A.R.T.-related functions

Detect Only:

When the S.M.A.R.T. function is enabled, the controller will send a command to enable all drives' S.M.A.R.T. function, if a drive predicts a problem, the controller will report the problem in an event log.

Detect and Perpetual Clone

When the S.M.A.R.T. function is enabled, the controller will send a command to enable all drives' S.M.A.R.T. function. If a drive predicts a problem, the controller will report the problem in an event log. The controller will clone the drive if a Dedicated/Global spare is available. The drive with predicted errors will not be taken off-line, and the clone drive will still act as a standby drive.

If the drive with predicted errors fails, the clone drive will take over immediately. If the problematic drive is still working and another drive in the same logical drive fails, the clone drive will resume the role of a standby spare and start to rebuild the failed drive immediately. This prevents a fatal drive error if yet another drive should fail.

Detect and Clone + Replace

The controller will enable all drives' S.M.A.R.T. function. If a drive predicts a problem, the controller will report the problem in the form of an event log. The controller will then clone the problematic drive to a standby spare and take the problematic drive offline as soon as the cloning process is completed.

Fail Drive

Before using this function, you should be ready with a hot spare so that a logical drive having a member disbanded can be quickly rebuilt. A disk drive can become unstable or dragging the array performance before being considered as a failed drive by your RAID system. If there are signs showing a member drive is seriously

degraded, (such as recurring reports of slow responses), you can use this option to disband a faulty drive from a logical drive once SMART-related errors are detected.

NOTE

The Fail Drive option can impose a danger in the situation when other members of a logical drive carry immanent defects. In the extreme cases, similar defects may be found in disk drives of the same lot by the same manufacturer. If you fail a member in a RAID5 array and another member encounters media errors during the rebuild process, you will lose data.

If you are using drives of different brands in your RAID system, as long as they are ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190-compatible, there should not be any problems working with the controller/subsystem.

You can use Galaxy Array Manager's Disk Performance Monitor to find out a low performing drive within a configured array.

14.3 Configuration Procedures

14.3.1 Enabling the S.M.A.R.T. Feature

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Periodic Drive Check Time

Follow the procedure below to enable S.M.A.R.T. on all drives.

1. First, enable the "Periodic Drive Check Time" function.

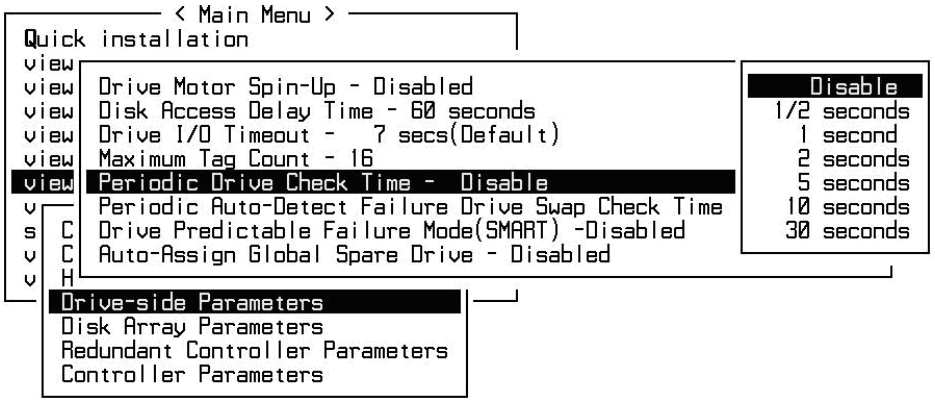

2. In "View and Edit Configuration Parameters" -> "Drive-side Parameters" -> "Drive Predictable Failure Mode <SMART>," choose one from "Detect Only,"

“Detect, Perpetual Clone” and “Detect, Clone+Replace.”

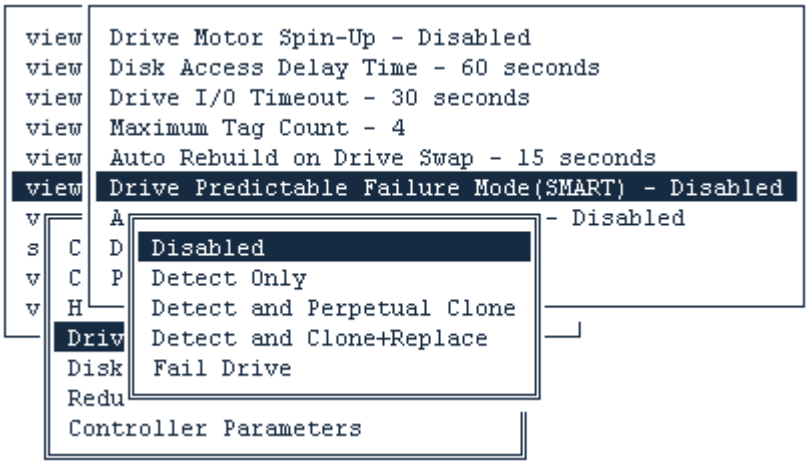
view Drive Motor Spin-Up - Disabled
view Disk Access Delay Time - 60 seconds
view Drive I/O Timeout - 30 seconds
view Maximum Tag Count - 4
view Auto Rebuild on Drive Swap - 15 seconds
view Drive Predictable Failure Mode(SMART) - Disabled
Disabled
Detect Only
Detect and Perpetual Clone
Detect and Clone+Replace
Fail Drive
Controller Parameters

14.3.2 Using S.M.A.R.T. Functions

Go to: View and Edit Configuration Parameters > Drive-Side Parameters > Drive Predictable Failure Mode (SMART)

Choose “Detect Only.”

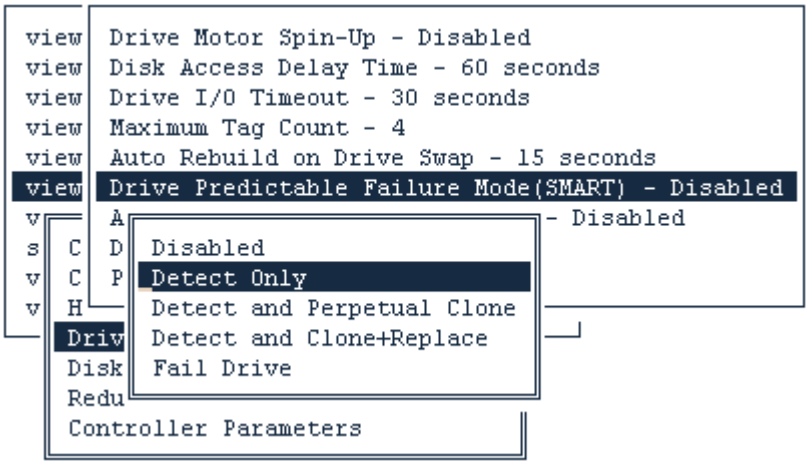
view Drive Motor Spin-Up - Disabled
view Disk Access Delay Time - 60 seconds
view Drive I/O Timeout - 30 seconds
view Maximum Tag Count - 4
view Auto Rebuild on Drive Swap - 15 seconds
view Drive Predictable Failure Mode(SMART) - Disabled
Disabled
Detect Only
Detect and Perpetual Clone
Detect and Clone+Replace
Fail Drive
Controller Parameters

Whenever a drive detects symptoms of predictable drive failure, the controller will issue an error message.

- The “Detect, Perpetual Clone” Setting

  Before selecting this option, you should make sure you have at least one spare drive for the logical drive (either Local Spare or Global Spare Drive).

  In “View and Edit Configuration Parameters” -> “Drive-side Parameters” -> “Drive Predictable Failure Mode <SMART>,” choose “Detect, Perpetual Clone.”

  When a drive (logical drive member) detects predictable drive errors, the controller will “clone” the drive with a spare drive. You may enter the "View and Edit Drives" menu and click on the spare drive (either Local or Global one). Choose from the menu items if you want to know the status of the source drive,

the cloning progress, or to abort cloning.

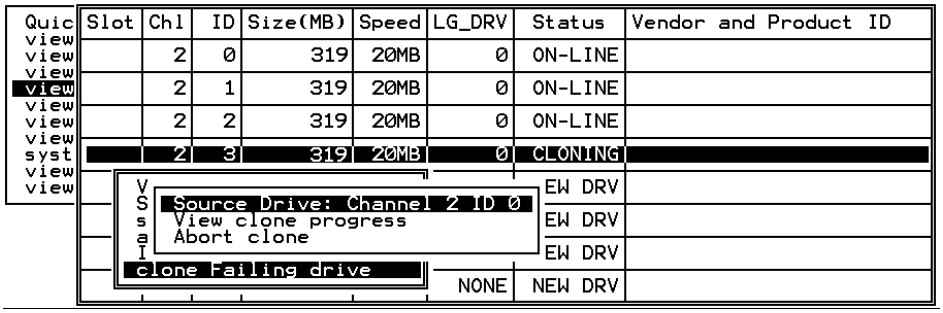

NOTE

As a precaution against the untimely failure of yet another drive, when configured as “perpetual clone,” the spare drive will only stay mirrored to the source drive (the drive with signs of failure), but not replace it until the source drive actually fails.

While the spare drive is mirroring the source drive, any occurrence of drive failure (when there are no other spare drives) will force the spare drive to give up the mirrored data and resume its original role – it will become a spare drive again and start rebuilding the failed drive.

- The “Detect, Clone + Replace” Function
  Before enabling this option, make sure you have at least one spare drive to the logical drive (either Local Spare Drive or Global Spare Drive).
  In “View and Edit Configuration Parameters” -> “Drive-side Parameters” -> “Drive Predictable Failure Mode <SMART>,” select “Detect, Clone+Replace.”
  When a drive (a logical drive member) detects the predictable drive failure, the controller will “clone” the drive with a spare drive. After the “clone” process is completed, it will replace the source drive immediately. The source drive will be identified as a “used drive.”
  If you want to see the progress of cloning, press [ESC] to clear the notification message and see the status bar.

The source drive’s status will be defined as a “used drive” and will be immediately replaced and pulled offline. This drive should be replaced with a new one as soon as possible.

15 Implementations for AV Applications

This chapter introduces new firmware functions that optimize array performance for AV applications.

> NOTE
>
> Due to the wide variety of I/O demands by different AV applications, detailed parameters such as read-ahead or cache threshold parameters can be otherwise implemented only by communicating with our technical support. This chapter only presents two generic configuration options. More options will be available for specific applications as customized features.
>
> All exemplary screens are captured from a hyper terminal management console.

15.1 AV Optimization Mode

The AV optimization option is applied for the emerging Audio/Video streaming applications such as the single-stream NLE (Non-Linear Editing), and the multi-stream VOD/MOD environments.

The AV Optimization Mode setting provides two configurable options: Fewer Streams and Multi-Streaming.

15.1.1 Fewer Streams: Read-ahead Performance

> Go to: View and Edit Configuration Parameters > Disk Array Parameters > AV Optimization Mode > Fewer Streaming

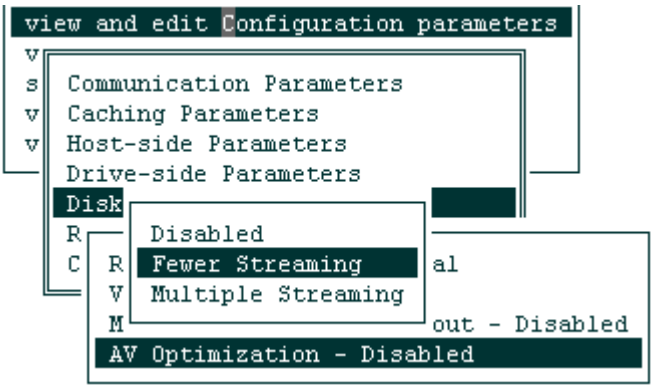

Applications such as an NLE (Non-Linear Editing) station may issue an I/O request for audio/video files of the sizes ranging from 1GB to 10GB or even larger.

Shown below is a RAID3 array configured in a 256KB stripe size. With only one 512KB outstanding I/O targeting at a large sequential file, the first I/O falls on two member drives while triggering a sequence of read-aheads at the same time. Read-aheads then occur across all member drives to make use of the combined disk performance.

The first I/O hit will be quickly returned and the read-aheads that ensue will be cached in memory. I/Os are then delivered through the read-aheads that are already stored in the fast data cache. As the result, applications featuring very few streams will be efficiently serviced through read-aheads in cache with minimized latency.

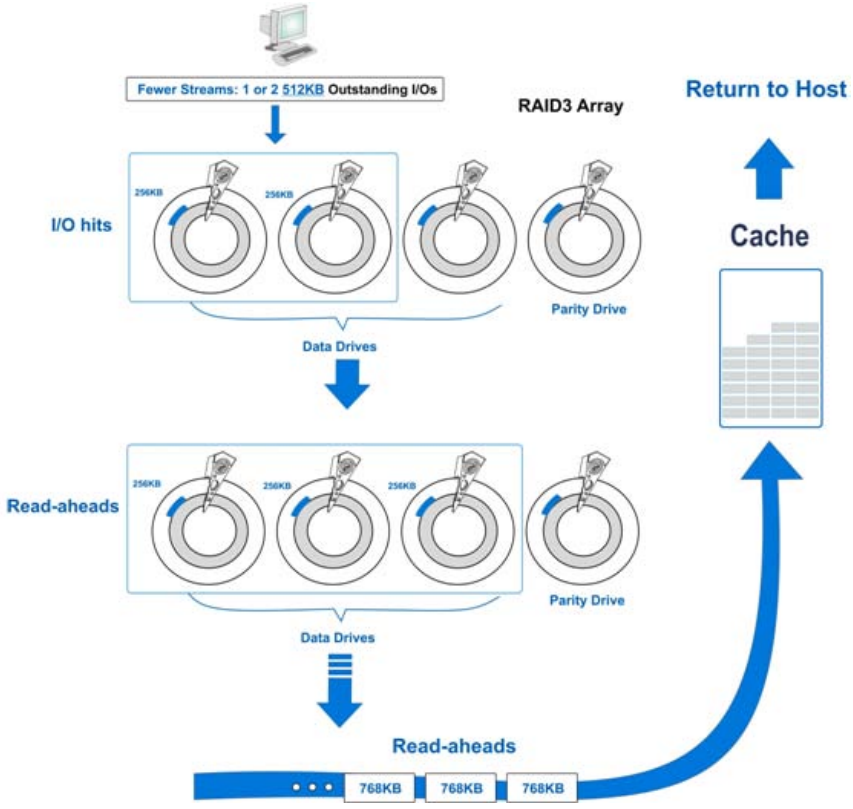

With the Fewer Streams setting, the related Maximum Drive Response Time is automatically set to 160ms to prevent interruptions by media errors.

15.1.2 Multi-Streaming: Simultaneous Access Performance

Go to: View and Edit Configuration Parameters > Disk Array Parameters > AV Optimization Mode > Multiple Streaming

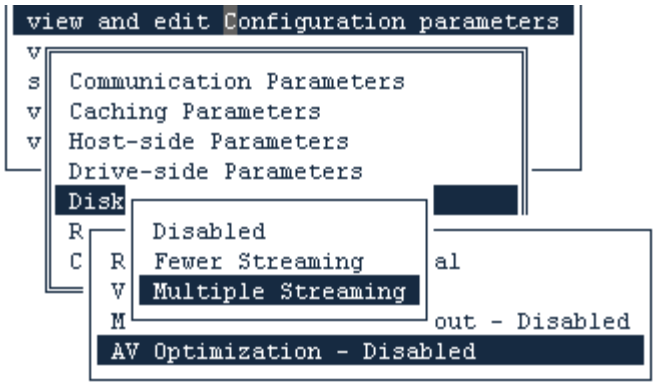

The Multi-Streaming option is designed for applications featuring shorter-length and concurrent requests coming in the swarm of outstanding I/Os, e.g., low-bit-rate clips in VOD or MOD Media Broadcasting.

Shown below is a RAID3 array configured in a 512KB stripe size. With multiple, say, 16 outstanding I/Os targeting at different data files, I/Os fall simultaneously on different member drives. As the result, each hard drive's actuator arms can quickly move to the next location to fulfill another I/O request.

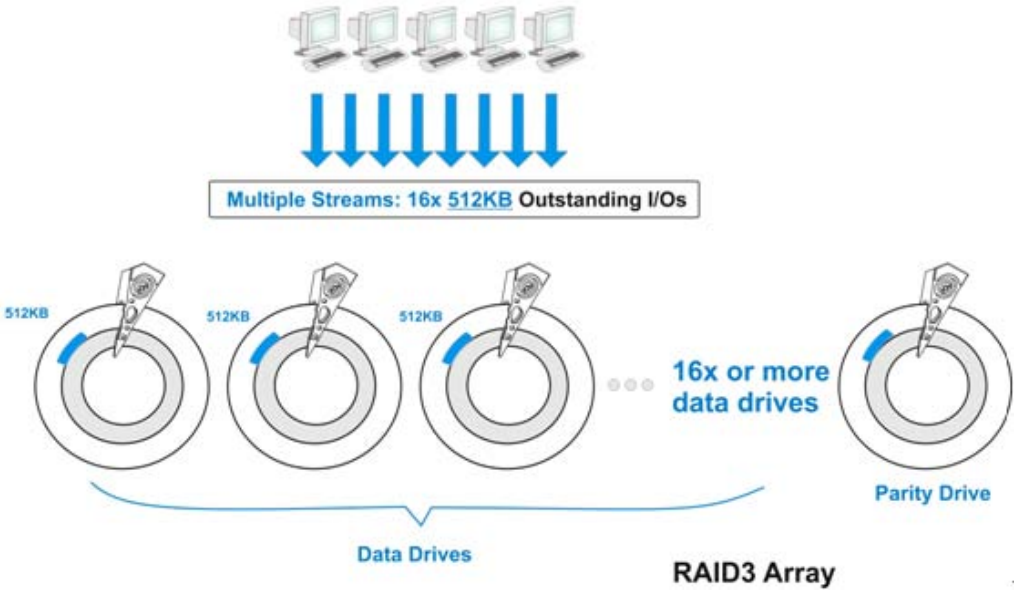

The Multi-Streaming option automatically configures the Maximum Drive Response

Time to 960ms because read latency cause less-serious problems with the smaller, randomly-generated file requests in VOD/MOD than the large, sequential files in NLE applications.

The Multi-Streaming applications require the following:

- A logical drive consisting sufficient number of disk drives so that I/Os can fall simultaneously on different members. Even though the real-world applications do not always make a perfect fit, configuring an array using an equal or slightly larger stripe size will ensure each individual outstanding I/O can fall within the range of a data drive's strip size (or chunk size).
- Properly tune the application I/O transfer size.
- Appropriate stripe size of your RAID arrays.

NOTE

The Maximum Drive Response Timeout bundled within the AV Optimization mode will over-rule any value you previously set in the similar menu found under Main Menu -> "View and Edit Configuration Parameters"-> "Disk Array Parameters."

15.2 Maximum Drive Response Time

Go to: View and Edit Configuration Parameters > Disk Array Parameters > Max Drive Response Timeout

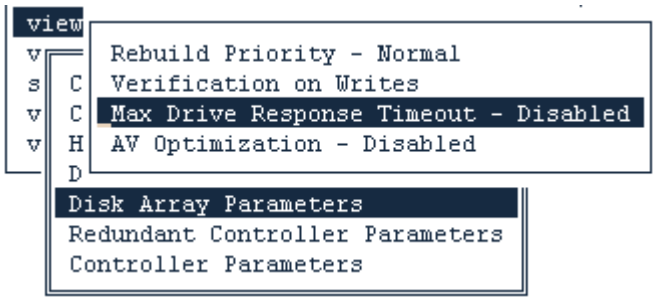

In situations such as drive failure or the occurrences of media errors, a read or write request returned after several hundreds milliseconds will be too long for AV applications for which choppy audio or dropped video frames are not acceptable.

15.2.1 Response Time in Read Scenarios

The Maximum Response Time option provides a timeout value for processing

read/write requests. If delays caused by media errors are reported on a specific member of an array, the subsystem firmware immediately retrieves data by generating data from RAID parity and the data blocks on other members of the array. In this way, delays on read requests can be efficiently eliminated. Without the Response Time setting, firmware may wait several seconds for the hard drive to timeout, which can be intolerable to some applications.

15.2.2 Maximum Drive Response Time in Write Scenarios

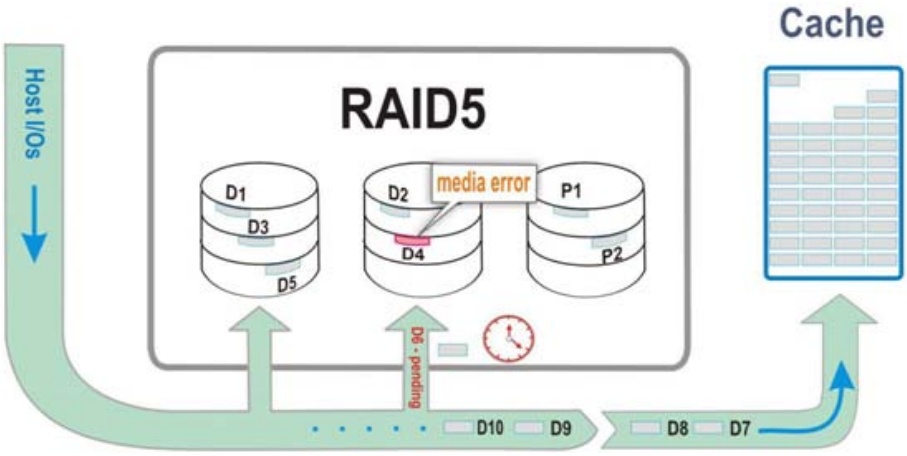

As shown above, the occurrences of media errors on a single disk drive can cause a performance drag within a few seconds. If media errors occur while servicing write requests, the following can occur:

- A media error is encountered while RAID system firmware is conducting a write request (D4: data block #4).
- It usually takes 3 to 4 seconds for a hard drive to return timeout state, and during that time the succeeding write requests (data blocks D7, D8, and onward) will be cached in system buffer and quickly fill the data cache.
- Supposed the data cache capacity is 512MB, it is easily used up when hundreds of megabytes of write requests come streaming down from the application server.
- When the cache is full, performance is quickly reduced and the benefit of write-back caching soon vanishes.

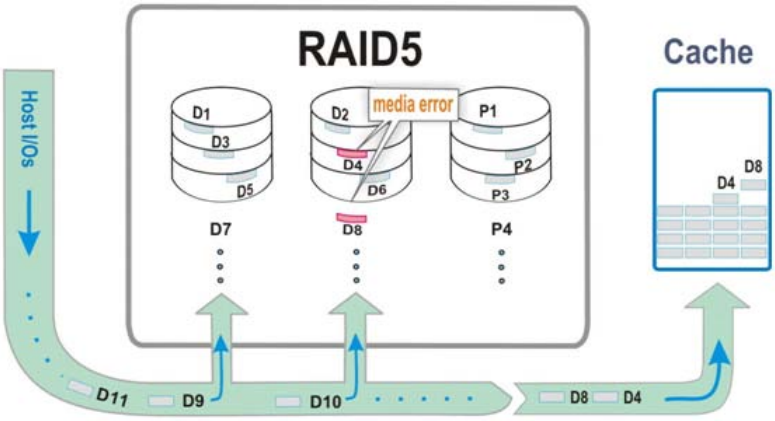

The Response Time remedy is described as follows:

- A response delay time is set in firmware: default is disabled.
- If a single disk drive cannot fulfill a write request within 160ms, firmware automatically proceeds with conducting write requests on other disk drives while also generating parity data.
- Only those writes affected by media errors on an individual disk drive will be cached in memory so that the data cache will not be quickly overwhelmed. The data cache holds a comparatively small portion of write requests. If a logical drive contains 8 members, one of them is parity drive and media errors are found on one member drive, caching data blocks to one disk drive only occupies 1/8 of cache capacity.
- With the response time on write, RAID subsystems can ensure array performance with the occurrences of media errors without waiting for physical hard drives to resolve hardware errors.
- If the drive carrying media errors does fail afterwards, data blocks cached in memory will be dumped and the rebuild begins.

15.2.3 Other Concerns

To prepare the array for read-intensive applications, the following are recommended:

- Default timeout as 160ms.
- Arrays should not be partitioned.
- The priorities for Rebuild or Media Scan operations should be set to “low.”
- Another timeout value, the “Drive I/O Timeout” which determines whether a drive has eventually failed to respond to I/O requests, is required as first-level timeout.

16 Redundant Controller

Sample topologies using redundant controllers can be found in the following discussions or in the Installation and Hardware Reference Manual that came with your RAID subsystems. The proceeding discussions will focus on the working theories and the configuration procedures for readying a redundant controller system.

16.1 Requirements

16.1.1 Concerns

Listed below are the configuration concerns and phenomena you will encounter when configuring a redundant controller subsystem:

- By system default, Controller A is always the primary RAID controller. Controller B in the lower slot serves as the secondary. If Controller A fails and is replaced afterward, firmware returns the Primary role to the replacement controller after a system reset.
- The traditional mapping method co-exists with the new, cross-controller access available with the latest firmware release.

NOTE

The latest HDX4 Firmware revisions support the Cross-controller ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs. However, mapping to both controllers' IDs is usually beneficial when it is difficult making the fault-tolerant links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. The Cross-controller mapping also makes sense in clustered server environments. Until now, SAS switch has not gained popularity on the market. For Fibre-host systems, fault-tolerant links can easily be made with the help of external bypass such as Fibre Channel switches.

For details of fault-tolerant link connections, please refer to your system Hardware Manual.

- One benefit of the cross-controller access is that when a host link fails, I/Os can travel through the counterpart controller, the RCC links, and then back to the RAID controller originally managing the array. The I/O load will still be managed by two controllers in the event of host link failure.

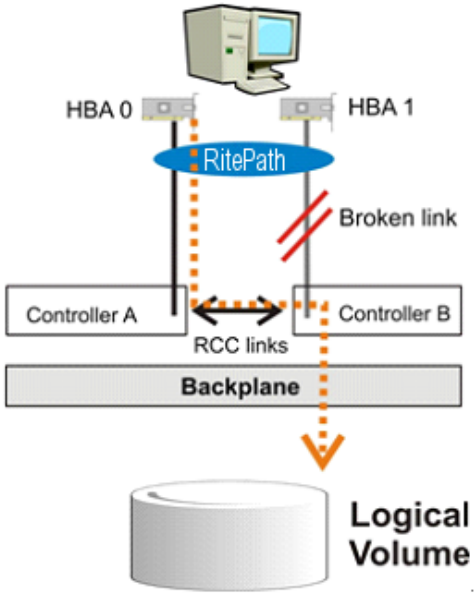

- If your subsystem comes with an LCD, the upper right corner of LCD will display a “P” or ”S,” meaning “Primary” or “Secondary” respectively. You may press the arrow keys together for two seconds to switch between the display of the Primary or Secondary controller status.
- The controller partners synchronize each other’s configurations at frequent intervals through the communications channel(s). And the synchronization act consumes part of the system resource.

16.1.2 Communications Channels

Controller Communications (Cache Synchronization) Paths:

Controller	RCC
Subsystem	Pre-configured RCC routes over the system backplane; may be SCSI, Fibre, or SATA data paths. These data paths cannot be re-assigned. The HDX4 series utilizes PCI-E channels for RCC traffic.
1U controller head	Older HDX: “Dedicated RCC” or “Drive+RCC.” Older HDX2: pre-configured RCC routes; no need to assign.

If controllers are running with write-back caching, a battery module on each controller is highly recommended.

16.1.3 Out-of-Band Configuration Access

RS-232C serial port cable (for terminal interface operation) connection. Normally a Y-cable will be included with dual-controller subsystems. The Y-cable ensures a valid link in the event of single controller failure.

Ethernet connection: If management through Ethernet is preferred, connect the Ethernet interface from both controllers to your local network. In the event of controller failure, the IP address assigned to the Primary Controller will be inherited by the surviving controller. In this way, the Ethernet port connection (management session) will be interrupted. An operator may have to re-enter the IP address to re-connect the controller/subsystem to a management console.

16.1.4 Limitations

- Both controllers must be exactly the same. Namely, they must operate with the same firmware version, the same size of cache memory, the same number/configuration of host and drive channels, etc. If battery backup is preferred, both should be equipped with a battery module.
- If a RAID controller fails and needs to be replaced, it is often the case that the replacement controller may carry a newer revision of firmware. It is advised you provide information such as firmware revision number, boot record version, etc. to your system vendor before acquiring for a replacement controller.
- For a subsystem featuring Fibre host channels and if the onboard hub is not enabled, connection through Fibre switches will be necessary for configuring fault-tolerant paths between host and RAID storage.
- In the event of data path failure, an intelligent FC switch should be able to direct data flow through an alternate path. In this case, multipathing software should be necessary to manage the data flow through the fault-tolerant paths that are strung between host and RAID storage.
- Your RAID subsystem may not come with sufficient numbers of Controller A and Controller B IDs. You will then need to manually create Controller A or Controller B IDs.

16.1.5 Configurable Parameters

Active-to-Active Configuration

Users can freely map a logical configuration to both the Controller A and Controller B IDs [putting forth different LUN views of a logical storage unit to different initiators (HBAs on servers)]. The I/O load to a logical drive can be dynamically shared by partner controllers.

The traditional mapping method requires at least two logical units which are separately managed by a RAID controller. Each logical unit is associated either with Controller A or Controller B IDs.

The dual-active configuration engages all system resources to performance. Users may also assign all logical configurations to one controller and let the other act as a standby (active-standby).

Cache Synchronization (Mirrored Cache)

The Write-back caching significantly enhances controller performance. However, if one controller fails in the redundant-controller configuration, data cached in its memory will be lost and data inconsistency will occur when the surviving controller takes over and attempts to complete the writes.

Cache synchronization distributes cached writes to both controllers and each controller stores an exact replica of the cache content on its counterpart. In the event of controller failure, the unfinished writes will be completed by the surviving controller.

16.2 Array Configuration Processes in Dual-controller Mode

16.2.1 General Firmware Configuration Procedures

Below are the basic procedures for readying a redundant-controller subsystem:

Step 1: Controller Unique Identifier

The Galaxy HDX4 subsystems usually come with a default identifier. If the default is lost for some reasons, provide a unique identifier for each RAID subsystem.

Go to: View and Edit Configuration Parameters > Controller Parameters > Controller

Unique Identifier

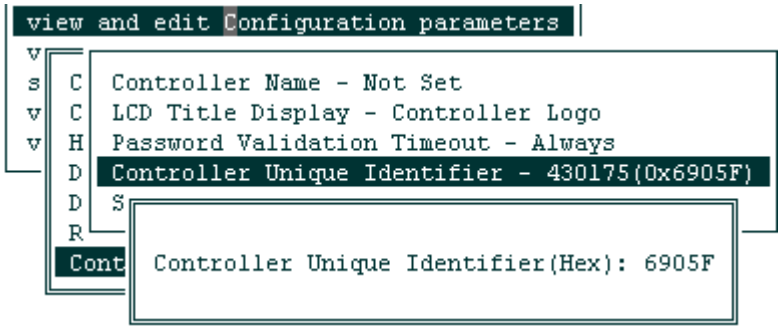
view and edit Configuration parameters
Controller Name - Not Set
LCD Title Display - Controller Logo
Password Validation Timeout - Always
Controller Unique Identifier - 430175(0x6905F)
Controller Unique Identifier(Hex): 6905F

Step 2: Create Controller A and Controller B IDs

1. "View and Edit Channels"☐ Choose a host channel.
2. "View and Edit ID"☐ Select an existing ID.
3. Under "Add/Delete Channel ID" ☐ "Controller A/Controller B" ☐ Select an ID from the pull-down list.
4. Reset the controller for the configuration to take effect.

Step 3: Create Logical Configurations of Drives

1. Under "View and Edit Logical Drives"☐ Select a Logical Drive entry.
2. Select a RAID level.
3. Select member drives
4. Configure other parameters, e.g., stripe size.
5. Create Logical Drives, Logical Volumes, and Logical Partitions..
6. Assign Logical Volumes either to the Slot A controller or to the Slot B controller.

Step 4: Map Each Logical Configuration of Drives to Controller A and/or Controller B IDs on host channel(s)

1. Under "View and Edit Host LUN"☐ Choose a "Channel-ID-Controller" Combination.
2. Select a “Logical Volume” and then the “Logical Partition” ☐ “Map to Host ID/LUN” (Create Host LUN Entry).
3. Repeat the process to avail a logical partition through multiple host IDs so that host can access the array through different host ports.

16.2.2 Setting Controller Unique ID (Optional)

> NOTE
>
> Each controller comes with a unique ID by default.

This value will be used to generate a controller-unique WWN node name, port names, Ethernet port MAC address, and to identify the controller during the failover process.

LCD Keypad

1. Galaxy HDX4 systems come with a default ID. It is recommended to use it. If the unique ID is accidentally lost, you can create a new ID using the following procedure:

```
View and Edit
Peripheral Dev..
```

2. Enter “View and Edit Config Parms”-> “Controller Parameters”. Use the up or down arrow keys to find “Ctlr Unique ID- xxxxx”.

```
Ctlr Unique (Hex)-
ID-
```

3. Enter a hex number from 0 to FFFFF and press [ENTER]. The value you enter should be different for each RAID subsystem.

Terminal

Go to: View and Edit Configuration Parameters > Controller Parameters > Controller Unique Identifier

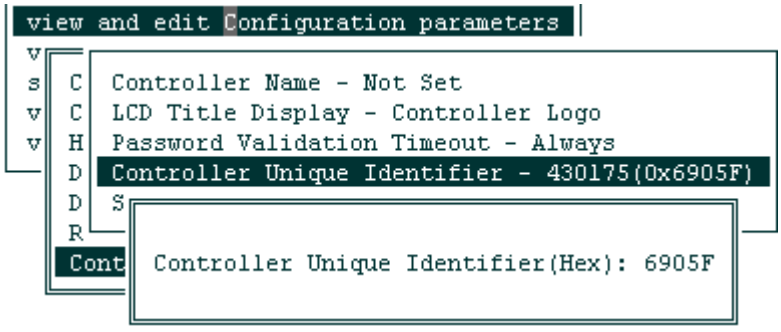
view and edit Configuration parameters
Controller Name - Not Set
LCD Title Display - Controller Logo
Password Validation Timeout - Always
Controller Unique Identifier - 430175(0x6905F)
Controller Unique Identifier(Hex): 6905F

16.2.3 Creating Controller A and Controller B IDs

The dual-controller operation may require that you manually create more Controller A and Controller B IDs.

LCD Keypad

1. In “View and Edit Channels”, press [ENT] to select a host channel.

```
CHL=0 ID=0
Slot A Ctlr  ..
```

2. Use the up or down arrow keys to select “Set Channel ID”. A pre-configured ID will appear, press [ENT] to proceed.

```
Add Channel
ID    ..
```

3. Use the up or down arrow keys to select “Add Channel ID” and then press [ENT] for two seconds on the “Slot A” or ”Slot B?” option to proceed.

```
Slot A
Controller    ?
```

4. When prompted by this message, use the arrow keys to select an ID. Press [ENT] to confirm.

```
Add CHL=0 ID=113
Slot A Ctlr  ?
```

5. A message will prompt to remind you to reset the controller. Press [ENT] to reset the controller or press [ESC] to return to the previous menu. The ID change will only take effect after a system reset.

```
Change Setting
Do Reset Ctlr ?
```

Terminal

Go to: View and Edit Channels > (Host Channel)

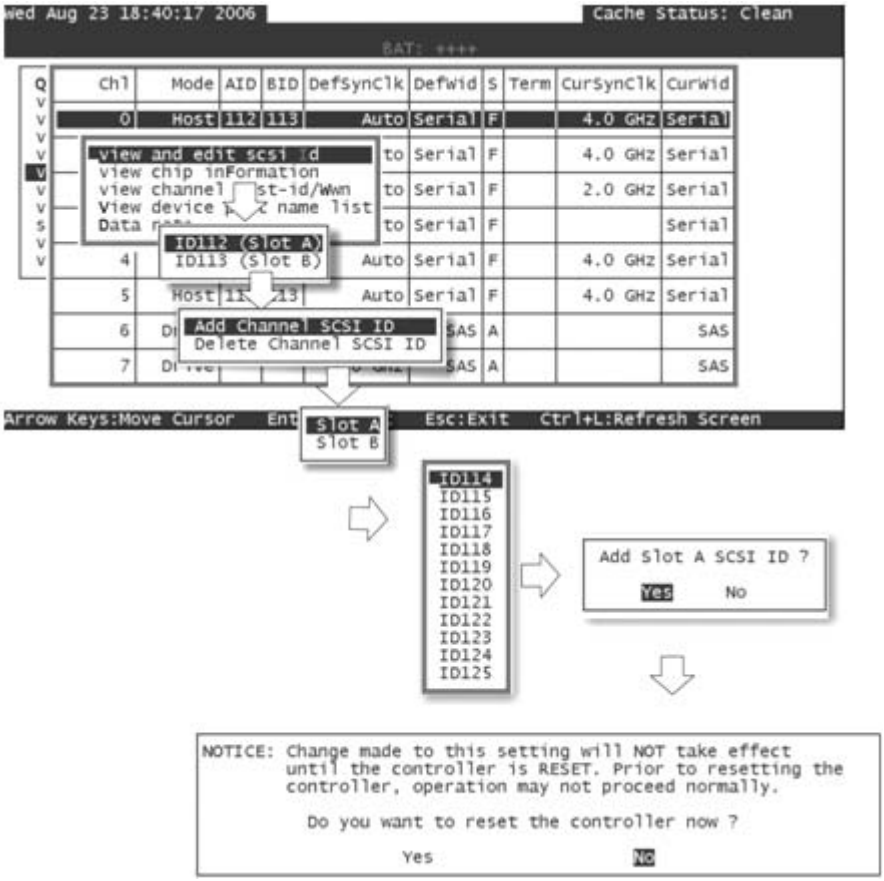
Wed Aug 23 18:40:17 2006
Cache Status: Clean
BAT: ++++
Ch1 Mode AID BID DefSyncClk DefWid S Term CurSyncClk CurWid
0 Host 112 113 Auto Serial F 4.0 GHz Serial
view and edit scsi id
view chip information
view channel host-id/WWN
View device port name list
Data rate
to Serial F 4.0 GHz Serial
to Serial F 2.0 GHz Serial
to Serial F Serial
ID112 (Slot A)
ID113 (Slot B)
4 Auto Serial F 4.0 GHz Serial
5 Host 11 13 Auto Serial F 4.0 GHz Serial
Add Channel SCSI ID
Delete Channel SCSI ID
6 SAS A SAS
7 SAS A SAS
Arrow Keys:Move Cursor Ent
Slot A
Slot B
Esc:Exit Ctrl+L:Refresh Screen
ID114
ID115
ID116
ID117
ID118
ID119
ID120
ID121
ID122
ID123
ID124
ID125
Add Slot A SCSI ID ?
Yes No
NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally.
Do you want to reset the controller now ?
Yes No

1. Select “Add Channel SCSI ID.” Press [ENTER] to confirm.
2. Select either “Slot A” or “Slot B” controller to create IDs that will be managed by a designated RAID controller.
3. A pull-down list will display all available IDs. Use your arrow keys to select an ID and press [ENTER] to confirm.

The configuration change will only takes effect after a system reboot.

16.2.4 Logical Volume Assignments (Dual-Controller Configuration)

A logical volume can be assigned either to Controller A or Controller B. By default, a logical volume is automatically assigned to Controller A, the controller installed in the upper slot (also the Primary controller by factory default). To divide the workload, you may manually assign a logical volume to Controller B.

NOTE

By default, logical drives and logical volumes will always be assigned to the Slot A controller. They can be manually assigned to the Slot B controller if the host

computer is also connected to the Slot B controller.

LCD Keypad

1. Press [ENT] key for two seconds to enter the firmware utility's Main Menu.
2. Use the arrow keys to navigate through the configuration menus. Choose "View and Edit Logical Volumes", then press [ENT].

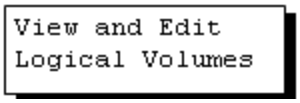

3. Create a logical drive or choose an existing logical drive, then press [ENT] to see the logical drive menu. The creation procedure is detailed in previous chapters.
4. Choose "Logical Volume Assignment..," then press [ENT].

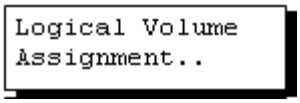

5. The message "Redud Ctlr LV Assign Slot B?" will appear. Press [ENT] for two seconds to confirm.

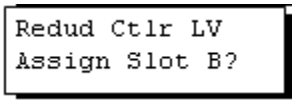

6. Map the logical partitions within logical volumes to a host ID or a LUN number under controller B ID. The host channel must have a "Slot B" ID. If not available, Slot B IDs can be manually added to a host channel.

Terminal

Go to: View and Edit Logical Volumes

1. Create a logical drive by selecting members and then a selection box will appear on the screen.

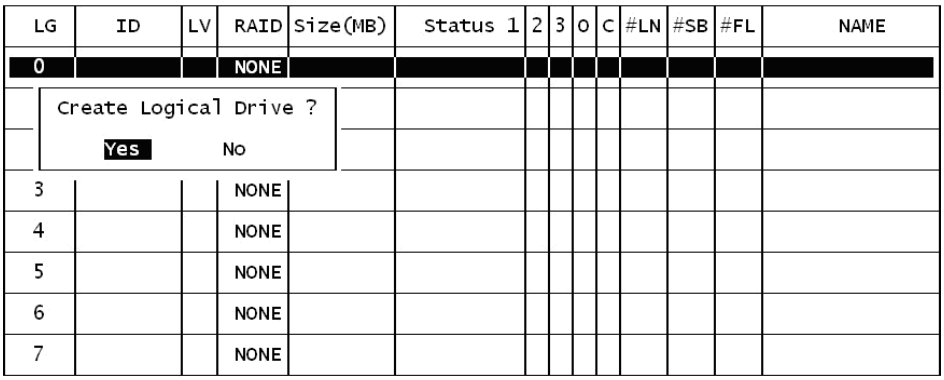

2. For the first logical drive on the RAID subsystem, simply select the first logical drive entry, LG 0, and press [ENTER] to proceed. You may create as many as 32 logical drives or more using drives in a RAID subsystem or in an expansion enclosure.

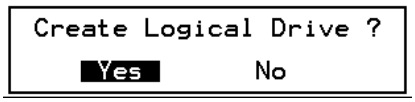

3. When prompted to “Create Logical Drive?,” select Yes and press [ENTER] to proceed. Please refer to the previous chapters for options specific to individual logical drives.

4. Select “View and Edit Logical Volumes” in the Main Menu to display the current logical volume configuration and status on the screen. Select a logical volume index number (0 to 7) that has not yet been defined, and then press [ENTER] to proceed.

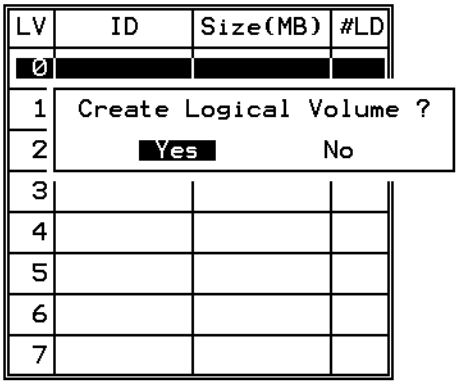

5. A prompt “Create Logical Volume?” will appear. Select Yes and press [ENTER].

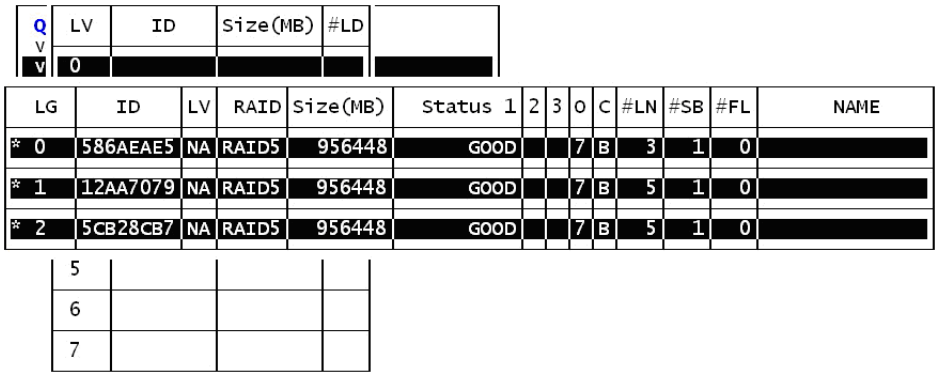

6. Select one or more logical drive(s) available on the list. An asterisk (*) will appear on the selected logical drive. Pressing [ENTER] again will deselect a logical drive.

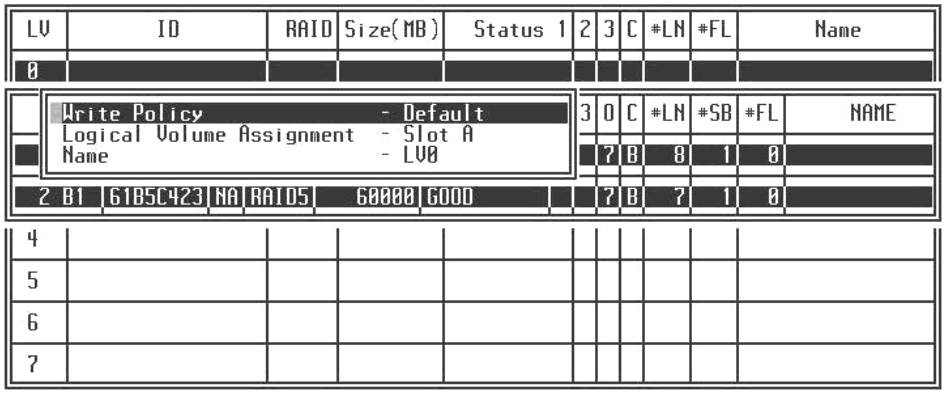

7. Use the arrow keys to select a sub-menu and change the write policy, controller assignment, and the name for the logical volume. You can balance the workload on partner RAID controllers by assigning volumes to both of them, e.g., 2 to Slot A controller and 2 volumes to the Slot B controller.

8. Logical volumes can be assigned to different controllers (primary or secondary; Slot A or Slot B controllers). The default is the primary or Slot A controller.

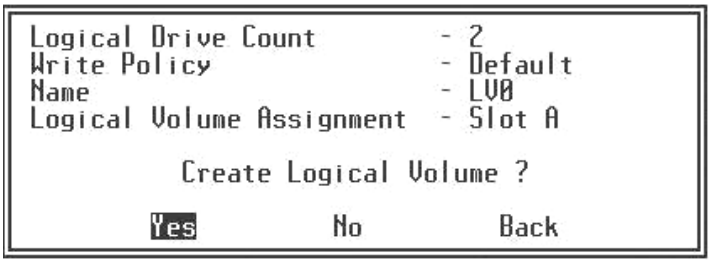

Note that if a logical volume is manually assigned to a specific controller, all its members' assignments will also be shifted to that controller.

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	0	C
A0	C54758	0	RAID5	140000	GOOD			?	B
A1	1E45E919	0	RAID5	120000	DRV FAILED			?	B
B2	182648C2	1	RAID5	50000	GOOD			?	B
B3	2AEB4D85	1	RAID5	50000	GOOD			?	B
4									

9. The reassignment is evident from the Logical Drive Status screen. "B2" indicates that the logical drive is Logical Drive #2 assigned to the Slot B controller.

> NOTE
>
> You cannot reassign a logical volume to a different controller until it is disassociated with host ID/LUNs (remove the previous LUN mapping).

16.2.5 Mapping a Logical Volume to Host LUNs

> NOTE
>
> Before proceeding with the mapping process, draw an abstract diagram of your configurations to help clarify the relationships among physical and logical components.
>
> Before the mapping process, check if you have properly configured logical drives, logical drive assignment, and host IDs. Changing host LUN mapping and re-configuring a RAID array may also require corresponding efforts on the management software running on host.

LCD Keypad

1. Choose "View and Edit Host Luns" from Main Menu and press [ENT] to proceed.

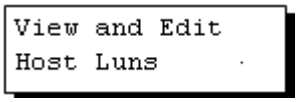

2. Use the arrow keys to navigate through the list of existing IDs and press [ENT] to select one of them.

```
Map Slot B
CH=0 ID= 000  ?
```

3. Use the arrow keys to choose a LUN number and press [ENT] to confirm.

4. Press [ENT] again to confirm.

```
Map Host LUN   ?
```

5. Use the arrow keys to select a logical volume if there are many.

```
LV0 RAID5 DRV=3
  9999MB GD SB=0
```

6. Press [ENT] and select a partition if the logical unit has been partitioned into individual capacity volumes.

```
LV=0 PART=0
     999MB   ?
```

7. Press [ENT] again to confirm.

```
CH0 ID9 LUN0 Map
to LV0 PRT0?
```

8. Press [ENT] to proceed.

```
Map Host LUN   ?
```

9. Press [ENT] to confirm.

```
Map Slot B
CH=0 ID= 0   ?
```

10. This message indicates that the logical unit has been successfully mapped to the ID/LUN combination. Use the arrow keys to continue mapping other logical units or press [ENT] to delete the mapped LUN.

```
CH0 ID9 LUN0
Mapto LV0 PRT0
```

11. Repeat the process to map all logical units to host ID/LUNs.

Go to: View and Edit Host LUNs > (Host)

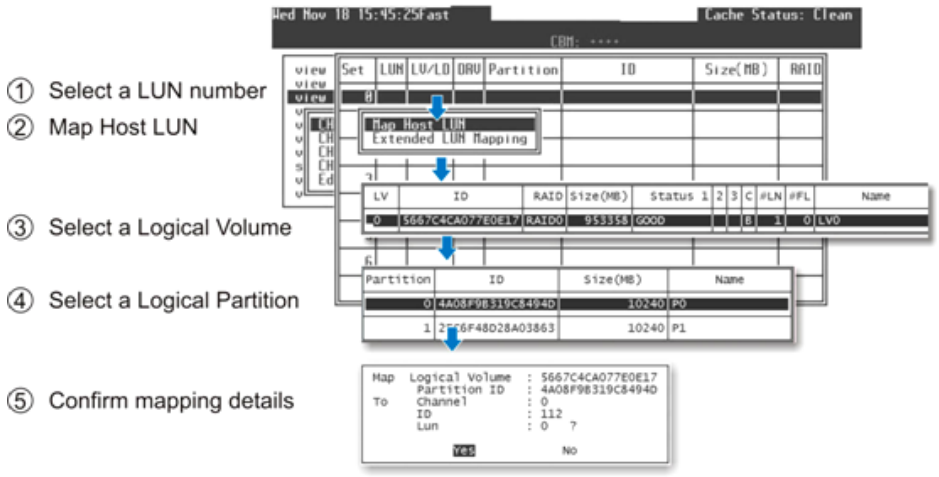

1. Select an LUN number under the host ID.
2. All logical volumes will be listed. Select one of them by pressing [ENTER] on it.
3. When selected, all logical partitions under the logical volume will be listed. Select a partition.
4. A confirm box will appear. Verify the details and press [ENTER] on Yes to complete the mapping process.
5. Repeat this process until you finish mapping all logical partitions to the host IDs you prefer. Repeat the process to map a logical unit to two host ID/LUNs if you want it to appear on two data paths.

16.3 Troubleshooting Controller Failure

16.3.1 What will happen when one of the controllers fails?

If one of the controllers fails, the surviving controller will automatically take over within a few seconds.

NOTE

Although the surviving controller will keep the system running, you should contact your system vendor for a replacement controller as soon as possible. Your vendor should be able to provide an appropriate replacement.

You should provide your vendor the serial number of the failed controller and the firmware version currently running on your system.

Some operating systems (e.g., SCO, UnixWare, and OpenServer) will not automatically retry with I/Os shortly delayed while the controller is taking over.

LCD Keypad

The red ATTEN LED on the LCD panel will light up, and the message "Redundant Ctlr Failure Detected" will appear on the LCD. Users will also be notified by audible alarm and messages sent over event notification methods such as Email, LAN broadcast, etc.

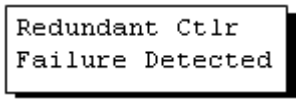

Terminal

1. When one controller fails, the other controller will take over in a few seconds.
2. There will be an alert message that reads "Redundant Controller Failure Detected."

3. Users will be notified by audible alarm and the messages through event notification methods such as Email, LAN broadcast, etc.
4. After a controller takes over, it will act as both controllers. If the Primary Controller fails, the Secondary Controller manages the logical drives originally managed by the Primary Controller.

16.3.2 When and how is the failed controller replaced?

Remove the failed controller AFTER the "working" controller has taken over. For the ventilation concern in RAID enclosures, it is better to leave a failed controller in place before a replacement arrives.

NOTE

If you need to replace a failed controller, DO IT WHEN THE SYSTEM IS POWERED ON AND IS MANAGED BY THE SURVIVING CONTROLLER! See the next section.

Redundant controller subsystems are designed to withstand a single controller

failure. If the replacement does not initialize properly, try the following:

When the replacement is connected, the "Auto-Failback" process should start automatically. If the replacement controller does not initialize, you may execute the following steps to bring the new controller online.

LCD Keypad

1. Press [ENT] for two seconds on the existing controller to enter the Main Menu.
2. Use the arrow keys to select "View and Edit Peripheral Dev..," then press [ENT].

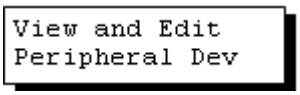

3. Choose "Set Peripheral Device Entry..," then press [ENT].

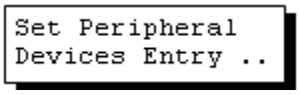

4. Select "Redundant Ctlr Function__," then press [ENT].

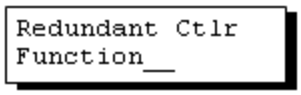

5. The message "Redundant Ctlr Slot A/Slot B Degraded" will appear on the LCD.

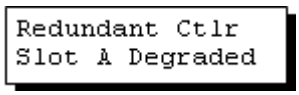

6. Press [ENT] and the message "Deassert Reset on Failed Ctlr?" will appear.

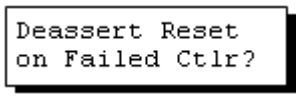

7. Press [ENT] for two seconds and the controller will start to scan for the new controller and bring it online.

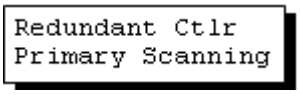

8. The new controller will then start to initialize.

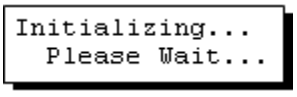

9. Once initialized, the replacement controller should assume the role of the Secondary Controller, and if the replacement is installed into the upper slot, it will restore its Primary role after a system reboot.

Terminal

Go to: View and Edit Peripheral Devices > Set Peripheral Device Entry > Redundant Controller > Force Primary Controller Failure

1. When the new controller is connected, the existing controller will automatically start initializing the replacement controller. If the replacement controller failed to initialize, try the following:

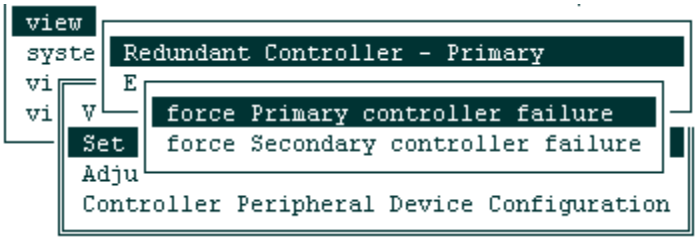

2. If the replacement has been initialized successfully, you may proceed to examine the system status. From the Main Menu, select "View and Edit Peripheral Devices" and then "View Peripheral Device Status" to see that the new controller is being scanned.

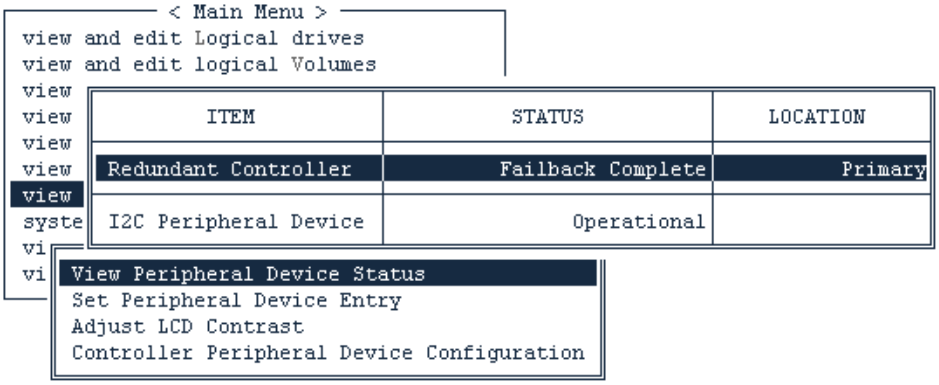

3. When the scanning is completed, the status will change to "Failback Complete."

16.3.3 How Do I Resolve Conflict in Assigning the Primary Controller?

Problems may occur if you replace a failed controller when system is powered down. If you power up both the surviving controller and the replacement together, they may contend for the role of the Primary (dominating) controller.

If you encounter this problem you may follow the procedure below to correct the fault:

LCD Keypad

1. 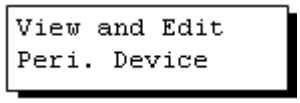

2. 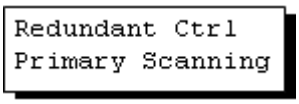

3. 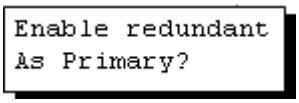

4. 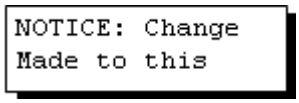

Go to: View and Edit Peripheral Devices > Set Peripheral Device Entry > Redundant Controller

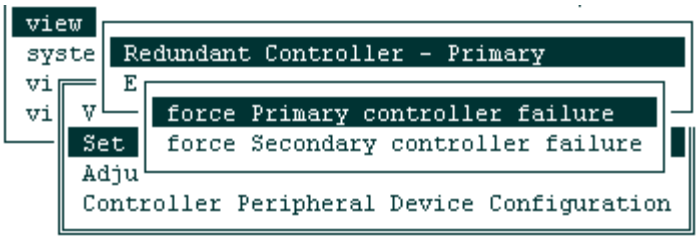

1. Stop host I/Os.
2. Power down the system and remove the surviving controller.
3. Power on and enter Main Menu -> View and Edit Peri. Device -> Set Peri. Device Entry -> “Redundant Controller” and change the controller role.
4. You may then install both controllers into their original positions and power on the RAID enclosure.

16.4 Configurable Parameters Related to Redundant Controllers

16.4.1 RCC (Redundant Controller Communications Channel) Status

Go to: View and Edit Configuration Parameters > Redundant Controller Parameters > Redundant Controller Communication Channel

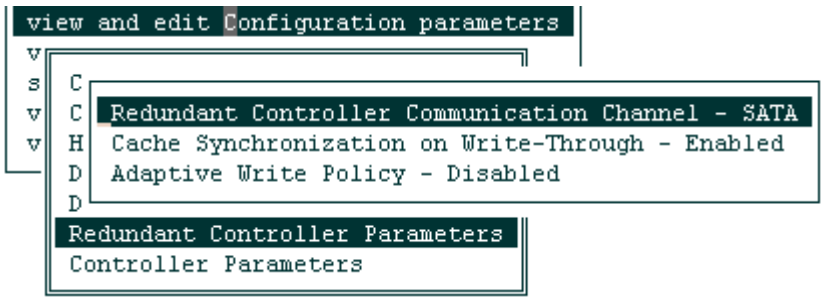
view and edit Configuration parameters
Redundant Controller Communication Channel - SATA
Cache Synchronization on Write-Through - Enabled
Adaptive Write Policy - Disabled
Redundant Controller Parameters
Controller Parameters

This item is for display only, showing the current communications routes.

16.4.2 Adaptive Write Policy

Go to: View and Edit Configuration Parameters > Redundant Controller Parameters > Adaptive Write Policy

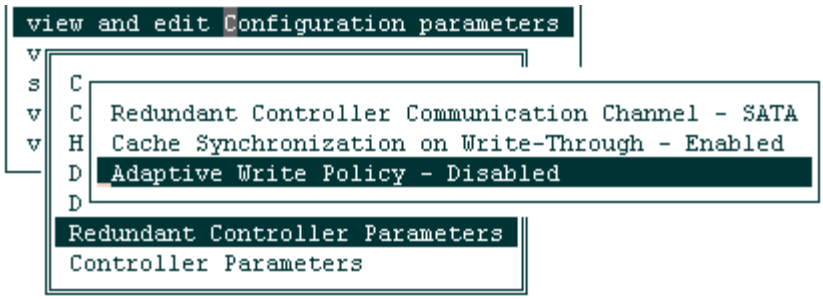
view and edit Configuration parameters
Redundant Controller Communication Channel - SATA
Cache Synchronization on Write-Through - Enabled
Adaptive Write Policy - Disabled
Redundant Controller Parameters
Controller Parameters

Firmware is embedded with intelligent algorithms to detect and to adapt the array's caching mode to the characteristics of I/O requests. The adaptive capability is described as follows:

1. When enabled, the Adaptive Write Policy optimizes array performance for sequential writes.
2. The adaptive policy temporarily disables an array's write-caching algorithm when handling sequential writes. Write-caching can be unnecessary with sequential writes for that write requests can be more efficiently fulfilled by distributing writes directly onto disk drives following the receiving order.
3. The adaptive policy changes the preset write policy of an array when handling I/Os with heterogeneous characteristics. If firmware determines it is receiving write requests that come in a sequential order, the write-caching algorithm is disabled on the target logical drives.

If the subsequent I/Os are fragmented and are received randomly, firmware automatically restores the original write-cache policy of the target logical drives.

16.4.3 Adaptation for the Redundant Controller Operation

If arrays managed by a redundant-controller configuration are configured to operate with write-back caching, cached data will be constantly synchronized between the partner controllers. Upon receiving sequential writes, firmware disables write-caching on target arrays and also the synchronized cache operation because the synchronization also consumes some of the controllers' processing power.

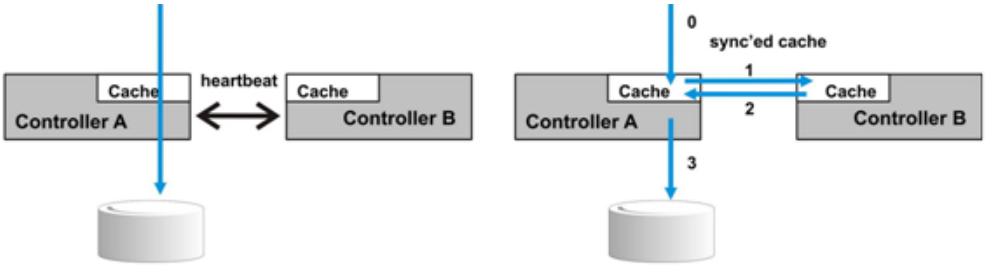

NOTE

Every time you change the Caching Parameters, you must reset the controller for the changes to take effect. The Adaptive Write Policy is applicable to subsystems working in the normal condition. If, for example, a drive fails in an array, firmware automatically restores the array's original write policy.

16.4.4 Cache Synchronization on Write-Through

Go to: View and Edit Configuration Parameters > Redundant Controller Parameters > Cache Synchronization on Write-Through

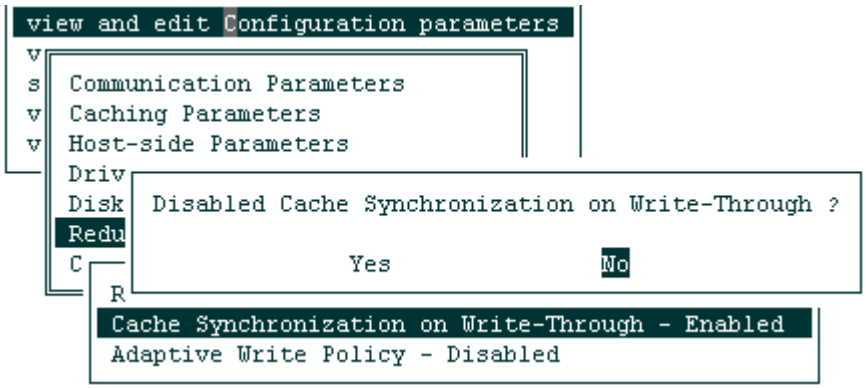

If your redundant controller system is not operating with Write-back caching, you can disable the synchronized cache communications between RAID controllers. Your system can be spared of the efforts to mirror and transfer data between partner controllers. This increases array performance for subsystems that operate without write caching.

Note that the configuration changes made to the RAID subsystem firmware will still be synchronized between the partner controllers.

16.5 Operation Theory

16.5.1 The Inter-Controller Relationship

The Primary/Secondary controller role is determined by a controller's position in a RAID enclosure. The new principle helps ensure the fixed location of a dominating, "Primary," controller. Other aspects of array management, ID/LUN mapping and array operation remain basically unchanged.

The new principle defines the RAID controller installed in Slot A, the upper controller slot, as the Primary controller. The factory configuration ensures that the Slot A controller always behaves as a Primary controller. In the following condition, a slot A controller temporarily serves as a Secondary controller:

1. If the Slot A controller fails, the original Slot B (Secondary) controller takes over and becomes the Primary controller.
2. When the slot A controller is replaced by a new controller, the new controller temporarily serves as the Secondary controller.
3. Once the subsystem is reset or powered down and powered on again, firmware returns the Primary role to the replacement controller in slot A.

16.5.2 Rules for Grouping Hard Drives and LUN Mapping

Listed below are the basics about configuring RAID arrays in a redundant-controller system:

- All configuration utilities are managed by the Primary RAID (normally the controller A) controller.
- Controller B status can also be displayed on a terminal or LCD screen. Management screen of a specific RAID controller is indicated by a flashing digit, <A> or <B> respectively on an LCD screen. Messages generated by different controllers will also be noted as shown below.

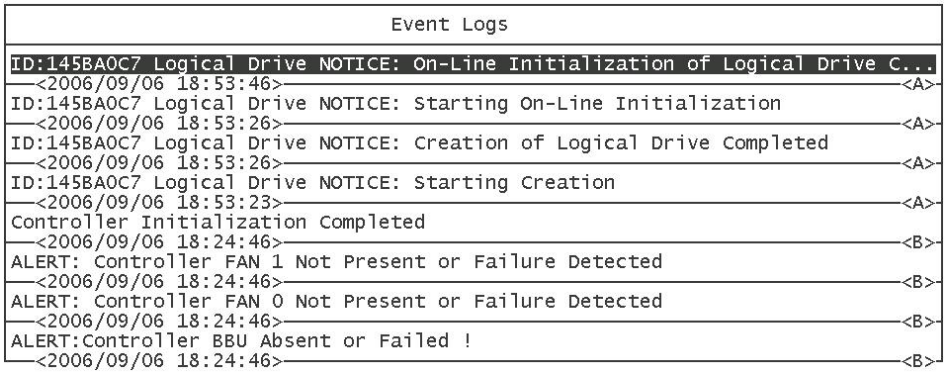

- In dual-controller mode, two controllers behave as one, and there is no need to repeat the configuration on another controller. The array configuration profile is automatically synchronized between partner controllers.
- Disk drive and array configuration processes are the same for subsystems using single or dual-active controllers.
- Using logical drives as the basic configuration units, system workload can be shared by partner RAID controllers. Logical units can be manually assigned to different controllers (Controller A or Controller B and consequently Primary or Secondary) to facilitate the active-active load-sharing configuration.
- Host channel IDs are designated either as a Controller A or as a Controller B ID. The controller A/B IDs then function as the designators for workload assigned to different RAID controllers.
- Each logical drive can be configured in a different RAID level.
- Several logical drives can be striped together to compose a larger logical volume. A logical volume then becomes the basic configuration unit for host LUN mapping and capacity management.
- Each of the logical units (a logical volume, or one of their partitions) can be made available on one or more host ports using the host LUN mapping function.
- Each of them can be “mapped” or “associated” with one or more host ID/LUNs. Each of these associated host ID/LUNs appears to the host operating system as a virtual storage volume.
- The idea is diagrammed as follows:

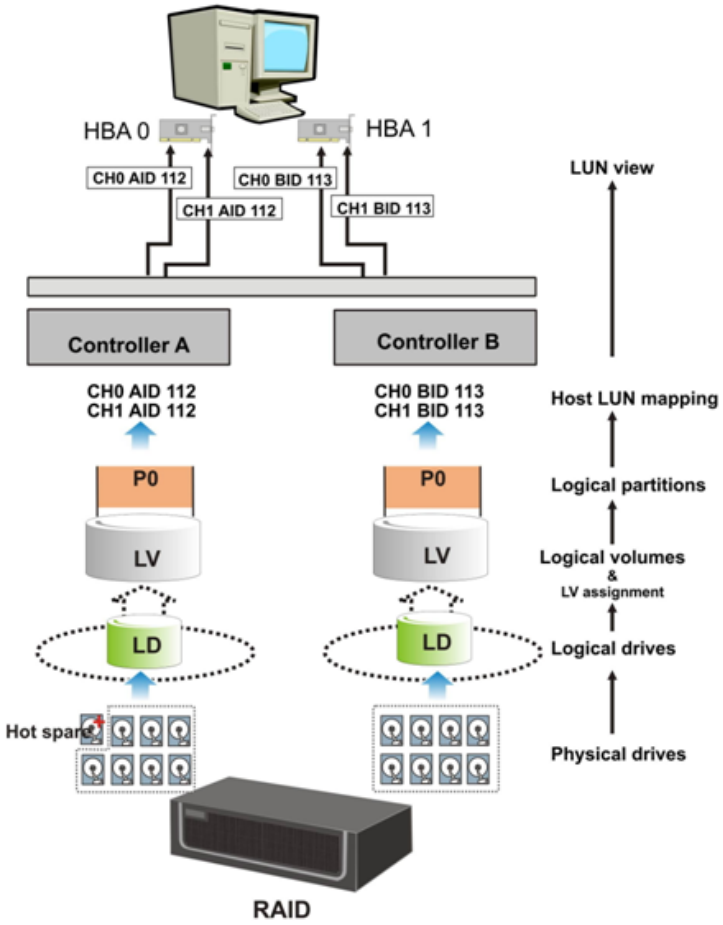

- As diagrammed below, array composition can be very flexible. You may divide a logical volume into several partitions, or use the entire logical volume as a single partition, with or without the support of spare drives.

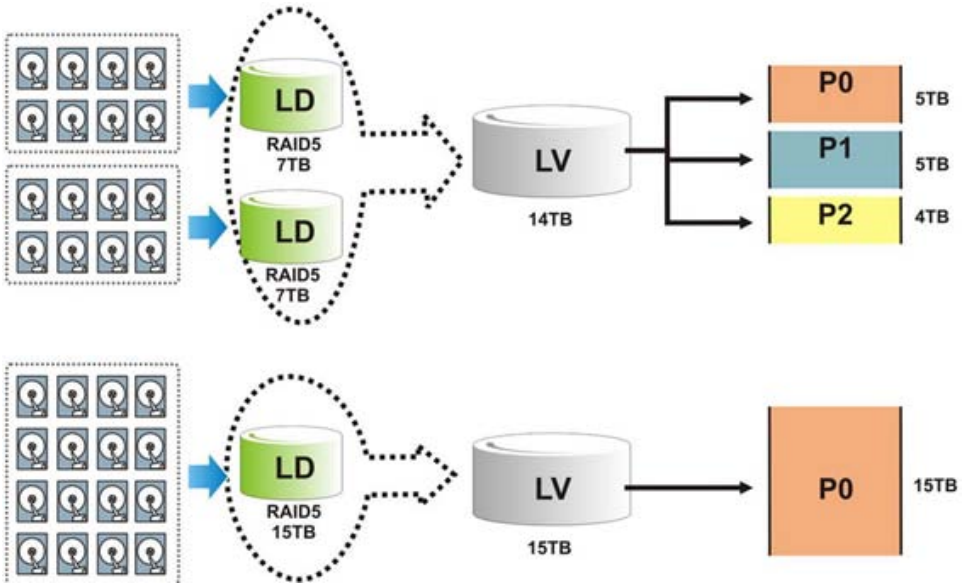

- Each logical partition can be associated (mapped) with one or more host IDs

(pre-configured as a Controller A or a Controller B ID) or the LUN numbers under these host IDs.

16.5.3 Host LUN Mapping: Design Concerns

- When it comes to building a reliable storage solution, redundancy is a virtue. We assume that an environment running mission-critical applications should consist of redundant RAID controllers and multi-pathing software that manage fault-tolerant data paths.
- Carefully configure your RAID arrays and select the appropriate settings such as stripe size and write policy. Reconfiguration takes time and may require you to move or back-up your data.
- Create at least two logical drives (LD0 and LD1) and associate (map) them equally with Controller A IDs (AID) and Controller B IDs (BID). Doing so you get the maximum work power from both of the RAID controllers.
- Logical RAID units are manually associated with Controller A or B IDs that reside on the host channels.
- Disable some configuration options for they might cause data inconsistency if module failures should occur. For example, disabling the use of buffers on individual disk drives may let you lose some performance, yet it is relatively safer for the drive buffers may hold cached writes during a power outage and cause data inconsistency.
- The configuration can be found in firmware's embedded utility through Main Menu -> View and Edit Configuration Parameters -> Drive-side Parameters -> Drive Delayed Write.

16.5.4 Mapping for Fault-tolerant Links

The purpose for mapping a logical drive to multiple IDs is diagrammed below:

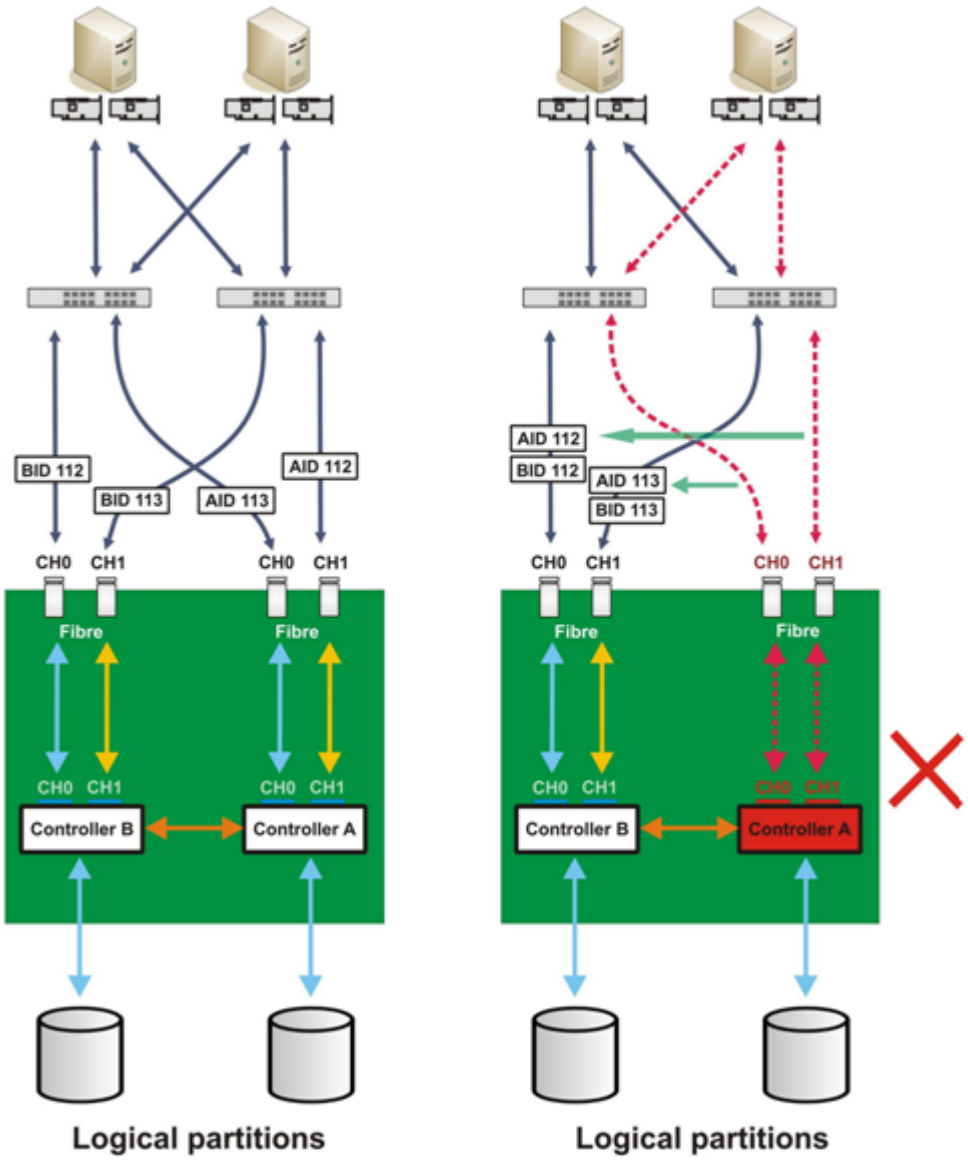

In the event of single RAID controller failure, all IDs managed by the failed controller will be taken over by the surviving controller. See the locations of mapped IDs in the above diagram.

If an application server can access the arrays through fault-tolerant paths, multi-path management software, such as Galaxy's RitePath, should be available.

Shown below is a condition with a broken host link. The host computer can still access the array (LD1) through an alternate data link. Even if one of the FC switches fails, access to data can still be continued:

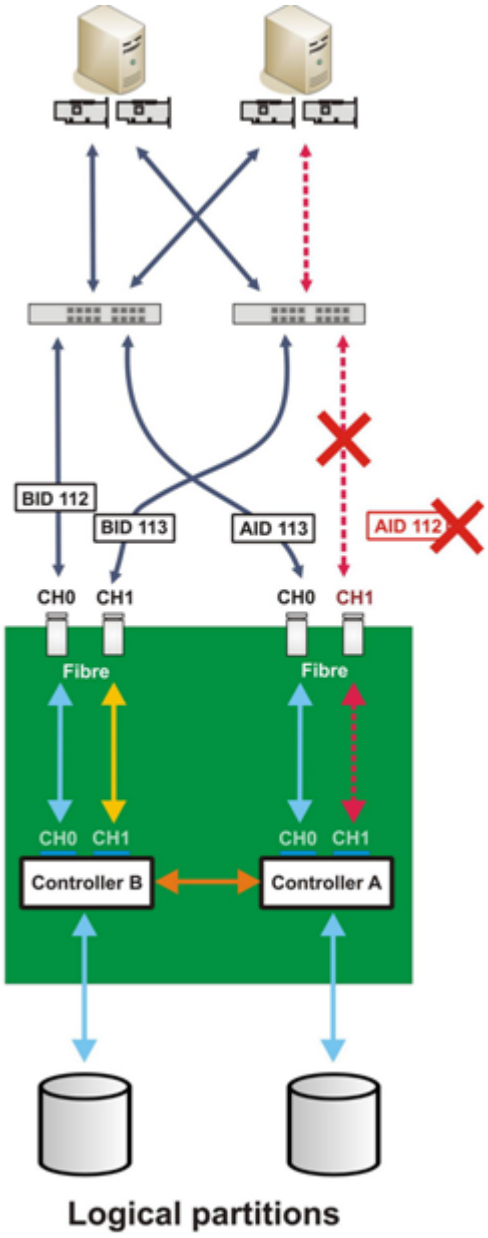
BID 112
BID 113
AID 113
AID 112
CH0
CH1
CH0
CH1
Fibre
Fibre
CH0
CH1
CH0
CH1
Controller B
Controller A
Logical partitions

16.5.5 Mapping Using the Cross-controller Mapping

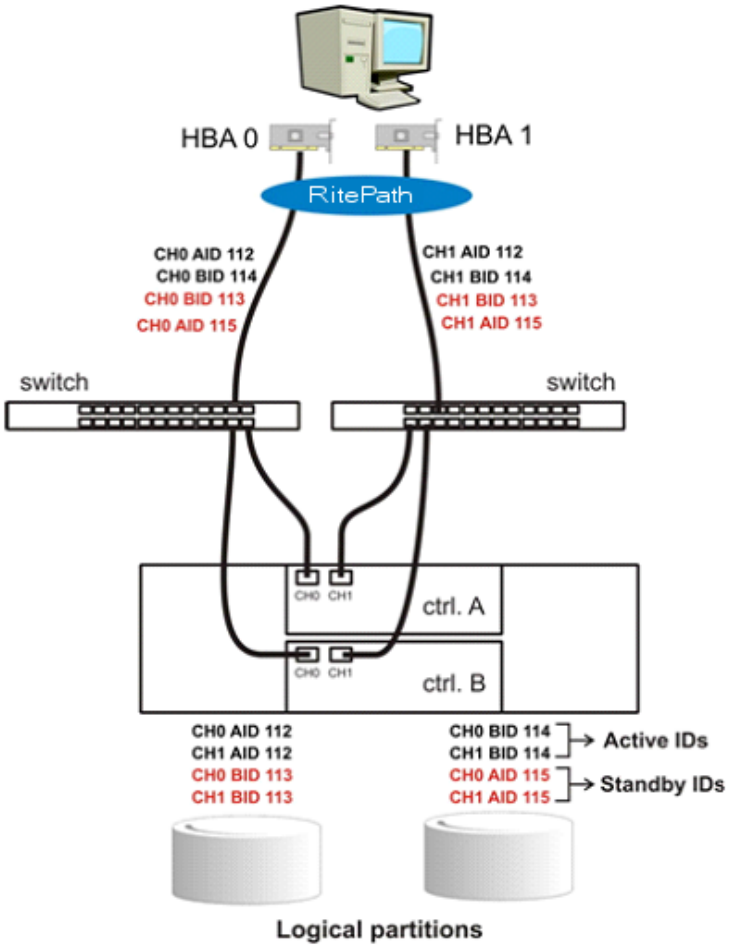

As diagrammed above, each logical partition is associated with two different channel IDs managed by different RAID controllers (AIDs or BIDs). This mapping method also ensures continuous host access in the situation when no port bypass is available, e.g., direct-attached without FC switches.

Note the following when configuring fault-tolerant configurations:

- Multi-pathing management software should be installed on the host computers to manage the access to the same array volume via two different I/O paths.
- Each channel ID (or an LUN under target ID) will appear as one virtual storage volume to the host operating system.

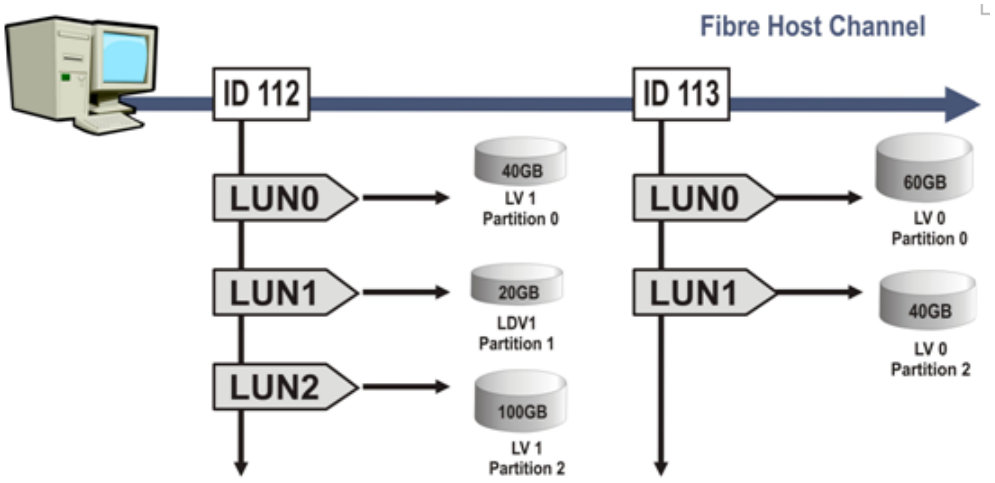

Shown below is a host channel bus teamed with multiple IDs/LUNs that are associated with logical partitions.

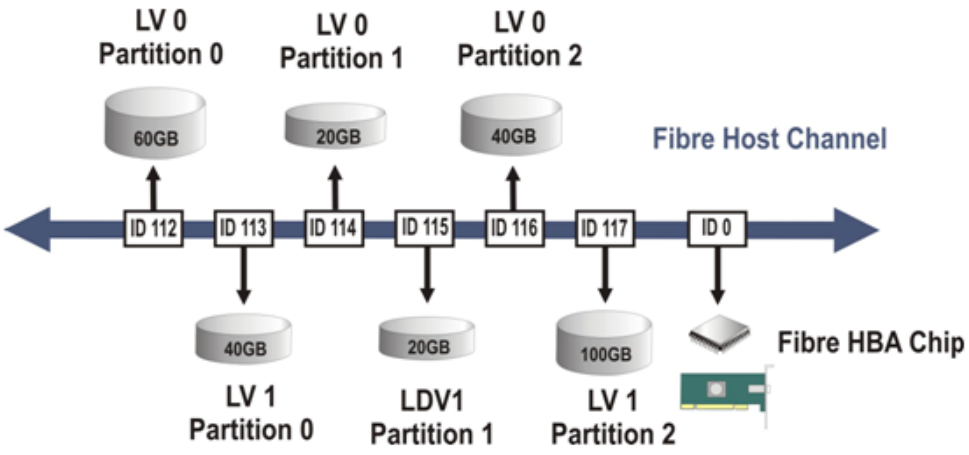

Some older operating systems/HBA cards do not read multiple LUNs under a target ID. As diagrammed above, you may have the host channel to present several IDs and map logical configurations to these IDs. Each of these IDs can be identified as “Controller A ID” or “Controller B ID.” As a rule for most operating systems, each configuration unit will then be mapped to LUN 0 under each ID.

16.5.6 Fault Tolerance

Why Use a Redundant Controller Configuration?

- Hardware failures can occur. A simple parity error can sometimes cause a RAID system to completely hang.
- Having two controllers working together will guarantee that at least one controller will survive catastrophes and keep the system working.

- When dealing with high-availability applications, redundancy is always a virtue. This is the logic behind having redundant controllers – to minimize the chance of down time for a storage subsystem.

A redundant-controller system uses two controller modules to manage the storage arrays. It requires two identical controllers to work together and both must be working normally. During normal operation, each controller serves its I/O requests. If one controller fails, the existing controller will temporarily take over for the failed controller. The failover and failback processes are completely transparent to the host (sometimes with the help of intelligent FC switch firmware) and require only minimum efforts to restore the original configuration.

Controller Failover and Failback

In an unlikely event of controller failure, the surviving controller will acknowledge the situation and disconnect from the failed controller. The surviving controller will then act as both controllers and serve all the I/O requests from host.

System failover is transparent to host. System vendors should be contacted for an immediate replacement of the failed unit.

Auto-Failback

Once the failed controller is removed and a replacement controller is installed, the existing controller will acknowledge the situation. The existing controller should automatically attempt to combine with the replacement controller.

When the initialization process of the replacement controller is completed, the replacement controller should always inherit the status of the Secondary controller.

NOTE

Reset the subsystem if the replaced controller resides in slot A. If the replacement controller in slot A is successfully combined, a system reset should restore its status as the Primary controller.

16.5.7 Fault Tolerance Procedures

1. Subsystem operating normally. Slot A controller is the Primary controller by factory default.

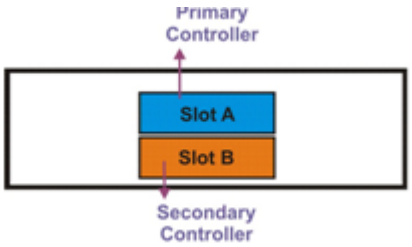

2. Slot A controller fails. Slot B controller inherits the Primary role.

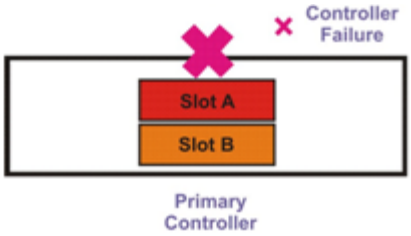

3. The failed controller in Slot A is replaced by a healthy replacement. The replacement controller becomes the Secondary controller temporarily.

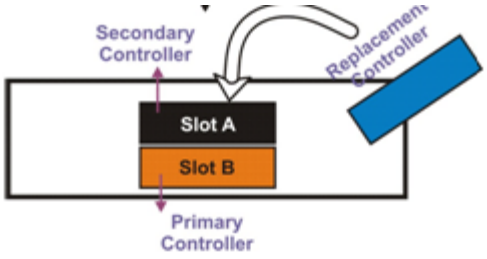

4. If the subsystem resets later, the Slot B controller returns the Primary role to the Slot A controller.

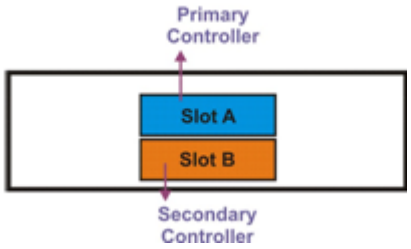

If the subsystem is reset later, the controller installed in the Slot A position will obtain the Primary controller status. The Slot B controller then resumes the Secondary role. The replacement controller will obtain all related configuration parameters from its counterpart.

16.5.8 Controller Failure

Controller failure is managed by the surviving controller (regardless of its original role as Primary or Secondary). The surviving controller disconnects from its counterpart

while gaining access to all signal paths. The existing controller then proceeds with the ensuing event notifications and take-over process.

Symptoms

- The LCD screen displays controller failure message.
- The surviving controller sounds an alarm.
- The "ATTEN" LED is flashing on the front panel.
- The surviving controller sends event messages to notify of controller failure (indicating its partner has failed).

16.6 Configuration Samples

16.6.1 Design Concerns

- We assume that an environment running mission-critical applications should consist of redundant RAID controllers and multi-pathing software that manages networking devices, such as FC switches or HBAs in fault-tolerant pairs.
- Carefully configure your RAID arrays and select the appropriate array settings such as stripe size and write policy. Reconfiguration takes time and may require you to move or back-up your data.
- Create at least two logical drives (LD0 and LD1) and associate (map) them equally with Controller A IDs (AID) and Controller B IDs (BID). Doing so you get the maximum work power from both of the RAID controllers. For more details on creating AIDs/BIDs and LUN mapping processes, please refer to the discussions later in this chapter.
- Logical RAID units are manually associated with Controller A or B IDs that reside on the host channels.
- Disable some configuration options for they might cause data inconsistency if module failures should occur. For example, disabling the use of buffers on individual disk drives may let you lose some performance, yet it is relatively safer for drive buffers may hold cached writes during a power outage and cause data inconsistency.

The configuration option can be found in firmware’s embedded utility through Main Menu -> View and Edit Configuration Parameters -> Drive-side Parameters -> Drive

Delayed Write.

There are similar concerns with the mirrored cache between the RAID controllers. Make sure compensatory measures are applied, e.g., use of battery backup modules or UPS devices.

Pros and Cons of Various Configurations

Configuration	Pros and Cons
Simple DAS w/o Hub	Applies to single logical drive over flexible cabling.
DAS w/ Hubbed Ports	DAS without FC switches; total host-side bandwidth can be halved by combining two host ports into a common host loop.
SAN w/ FC Switches	Applies to multi-server SAN; requires external FC switches.
Multi-pathing w/ Clustered Servers	High redundancy on server side and on the storage side. I/O path re-routing is partially managed by FC switches.

16.6.2 Simple DAS without Hub (Cross-controller Mapping Method)

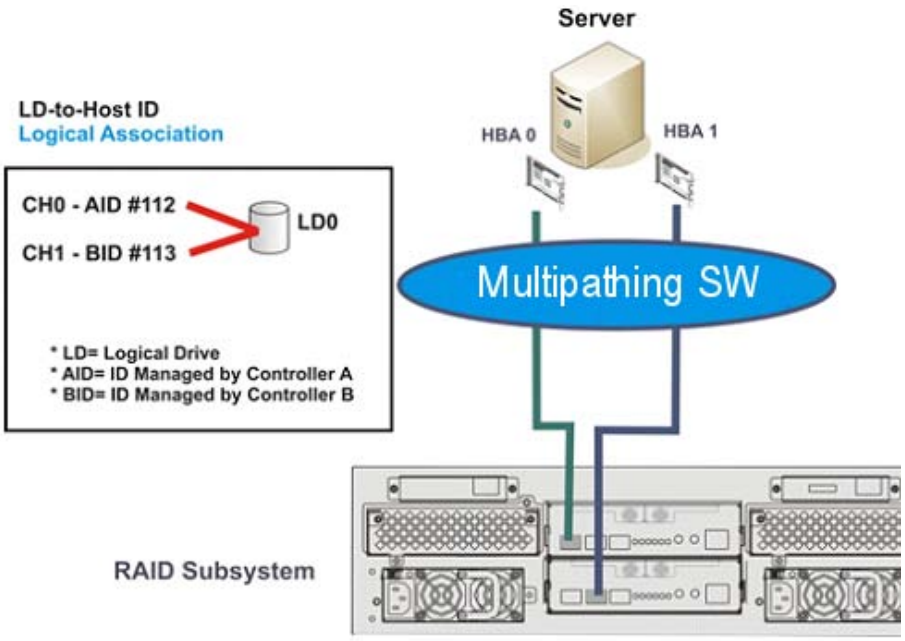

Tasks	Logical Drive	LUNs	Channel	AID	BID
Map LD0 to an AID on channel #0.	LD0	0	0	112	N/A
Map LD0 to a BID on channel #1 for redundant-path access.	LD0	0	1	N/A	113

This configuration applies to a dual-controller subsystem directly attached to a host computer without intermediate networking devices. A logical drive is associated with different Controller IDs (Controller A and Controller B IDs) on separate host channels and different RAID controllers.

In the event of cabling or controller failure, host can still access the array.

NOTE

You may use different channel IDs than are shown here in the sample topologies, IDs used in the sample configurations are mostly default numbers in firmware. As long as the IDs are carefully selected according to the configuration rules, there is no limitation on using different host channel IDs.

A logical drive is associated with both a Controller A and a Controller B ID. This methodology applies when no onboard or external bypass is available.

You may use the onboard hub to combine two host ports into a common host loop. Then you may not need the cross-controller mapping.

16.6.3 SAN with FC Switches

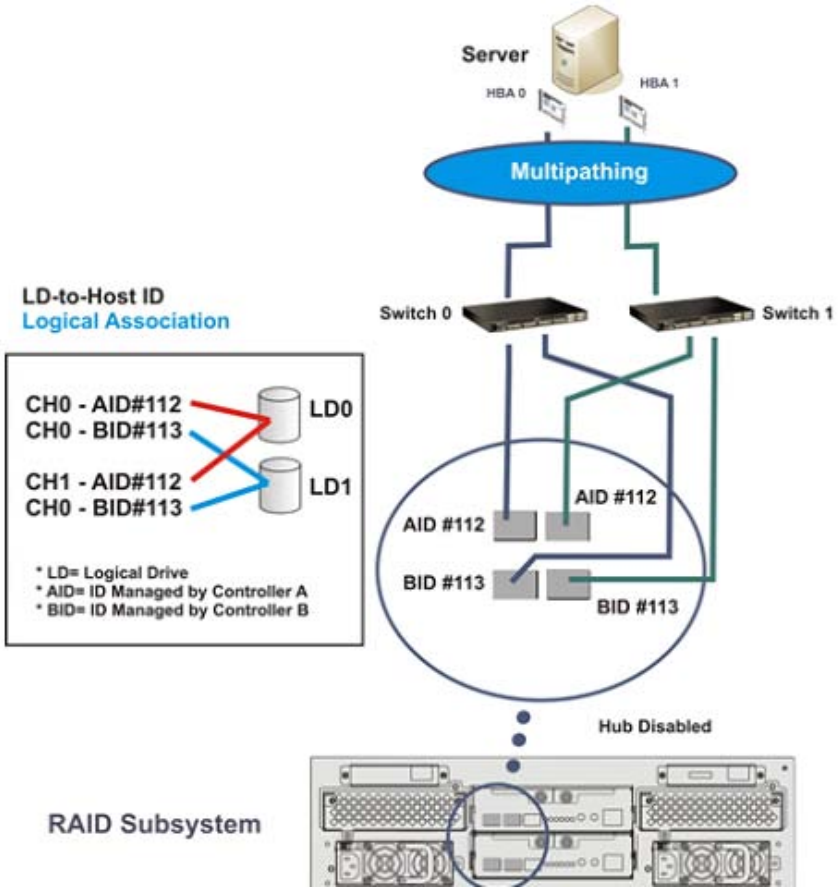

Shown above is a configuration using FC switches to facilitate the connections with multiple SAN servers. For the reason with diagram's simplicity, only one server is displayed.

Tasks	Logical Drive	LUNs	Channel	AID	BID
Map LD0 to an AID on channel #0.	LD0	0	0	112	N/A
Map LD0 to an AID on channel #1 for path redundancy.	LD0	0	1	113	N/A
Map LD1 to a BID on channel #0.	LD1	0	0	N/A	113
Map LD1 to a BID on channel #1 for path redundancy.	LD1	0	1	N/A	112

This configuration applies to a redundant-controller subsystem attached to switched fabric and then to application server(s).

Fault Tolerance is achieved through the following:

- Logical drives are separately associated either with the Controller A IDs or Controller B IDs on separate host channels.
- In the event of a controller failure, the surviving controller inherits IDs from the failed controller. Host IDs managed by a failed controller are automatically passed down to a surviving RAID controller. For instance, Controller A IDs will be managed by the Controller B if Controller A fails.
- In the event of cabling failure, an array is access through the alternate data path through an alternate host ID.
- Through the intermediate FC switches or switch zoning, cable/controller failure can be managed by re-routing host I/Os to a valid link.
- When attached to switched fabrics, the subsystem’s onboard hub function should be disabled.

16.6.4 Multi-pathing with Clustered Servers (Cross-controller Mapping Method)

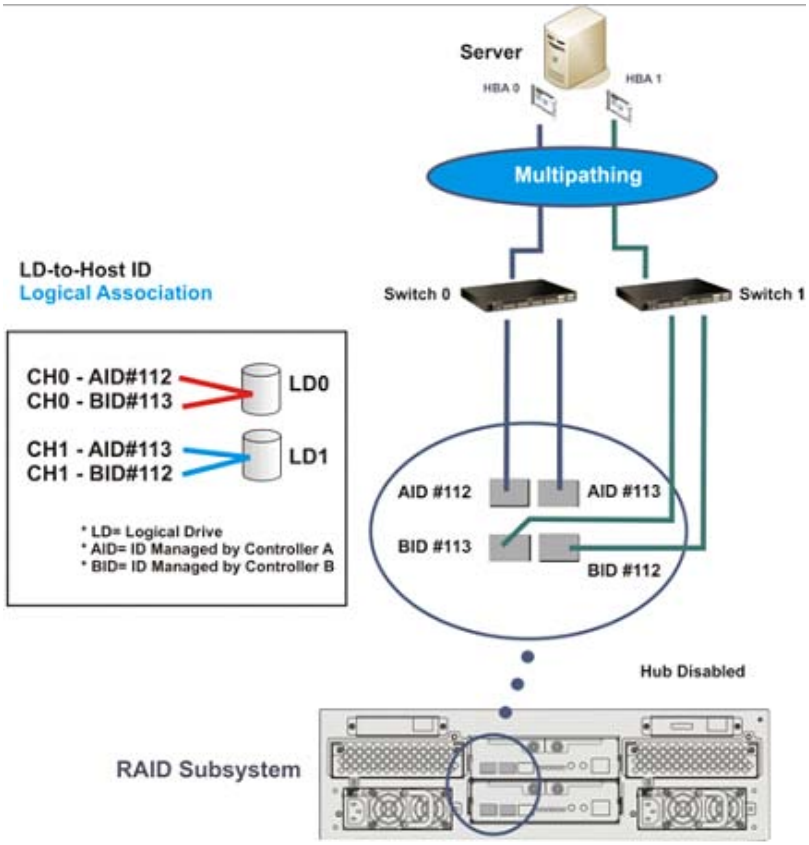

Tasks	Logical Drive	LUNs	Channel	AID	BID
Map LD0 to an AID on channel #0.	LD0	0	0	112	N/A
Map LD0 to an AID on channel #1 for path redundancy.	LD0	0	0	N/A	113
Map LD1 to a BID on channel #0.	LD1	0	1	113	N/A
Map LD1 to a BID on channel #1 for path redundancy.	LD1	0	1	N/A	112

The multi-pathing software is installed on both of the clustered servers to manage the fault-tolerant data paths

17 Firmware Functionality Specifications

17.1 Basic RAID Management

RAID levels	0, 1(0+1), 3, 5, 6, 10, 30, 50, 60, JBOD and NRAID. Levels 10, 30, 50, and 60 are the multi-level RAID defined as the logical volume implementations; logical volumes consist of logical drives of different RAID levels that are striped together. Including logical drives of different RAID levels in a logical volume is, however, not recommended.
Maximum number of logical drives	up to 64 with a 1GB or above memory size
Maximum logical drive capacity	64TB
RAID level dependency to each logical drive	Independent. Logical drives configured in different RAID levels can co-exist in a logical volume and within a RAID subsystem
Maximum number of logical drive members	128 with 512MB memory size (specification number, not recommended for the difficulties with backup, rebuild, and management tasks)
Configurable stripe size	16KB, 32KB, 64KB, 128KB, 256KB, 512KB, or 1024KB per logical drive
Configurable Write Policy (write policy per array)	Write-Back or Write-Through per logical drive. This policy can be modified later.
Logical drive identification	Unique, controller randomly generated logical drive ID; Logical Drive and Logical Volume name user-configurable for ease of identification in a multi-array configuration
Maximum number of partitions for each logical drive	up to 64 with a 1GB memory size
Maximum number of logical volumes	16 with a 1GB or above memory size
Maximum number of LUNs Mappable	up to 1024 with a 1GB or above memory size

Maximum number of LUNs per host ID	Up to 32, user configurable
Maximum number of Media Scan task schedules	16
Concurrent I/O	Supported
Tag Command Queuing (TCQ)	Supported
Native Command Queuing (NCQ)	Supported
Dedicated spare drive	Supported, hereby defined as the spare drive specifically assigned to a logical drive. Also known as Local Spare
Global spare drive	Supported, the spare drive that serves all logical drives (as long as it is equal in size or larger than logical drive members)
Global spare auto-assign	Supported, applies to all unused drive(s); safeguards the array if a spare has been used in the previous array rebuild and users forget to configure a new drive as a spare.
Enclosure spare drive	A Spare that participates only in the rebuild of a failed drive within the same enclosure.
Co-existing Dedicated (Local), Enclosure-specific, and Global spare drives	Supported
Auto-rebuild onto spare drive	Supported
Auto-scan of replacement drive upon manually initiated rebuild	Supported
One-step rebuild onto a replacement drive	Supported
Immediate logical drive availability	Supported; Logical arrays are immediately ready for Host I/Os. Initialization task is completed in the background except when the logical array is stated as “INCOMPLETE” or “BAD;” e.g., has a failed member right after the creation.
Auto-rebuild onto failed drive replacement	Supported. With no spare drive, the subsystem will auto-scan the failed drive and starts rebuild automatically once the failed drive has been replaced.
Concurrent rebuild / expansion	Multiple logical drives can proceed with a Rebuild/Regenerating Parity, and/or Expansion/Initialization/Add Drive operation at the same

	time. NOTE: Regenerate Parity and Rebuild cannot take place on a logical drive at the same time. Create, Expand, and Add Drive operations cannot take place on a logical drive at the same time.
Background firmware download	Firmware can be downloaded during active I/Os, and takes effect after a system reboot.
Auto recovery from logical drive failure (configuration on drives)	Supported. If a user accidentally removed the wrong drive to cause the 2nd drive failure of a one-drive-failed RAID5 / RAID3 logical drive, fatal error may occur. However, you may force the system to reaccept the logical drive by switching off the subsystem, installing the drive back to its original drive slot, and then power on the subsystem. You may have the chance to restore the logical drive into the one-drive-failed status.

NOTE:

To ensure smooth operation, sufficient cache memory buffer is required for configurations made up of numerous logical units. An intelligent trigger mechanism is implemented in the latest firmware version 3.85 and later. If a subsystem/controller comes with a DIMM module of the size equal or larger than 1GB, firmware automatically enlarges the maximum numbers of logical units.

	DIMM size < 1G	DIMM size >= 1G
Max. no. of LD	16	32
Max. no. of LV	8	16
Max. partitions per LV	16	64
Max. no. of LUN	128	1024

17.2 Advanced Features

Media Scan	Supported. Verify written data on drives to avoid bad blocks from causing data inconsistency. If bad blocks are found, data can be reconstructed by comparing and recalculating parity from adjacent drives (RAID1/3/5/6). The “Reconstruction Writes” are followed by “Write Verification” operation.

Bad Block Handling in degraded mode	A method for handling low quality drives. The operation is performed on both the logical drive in degraded mode or those that are being rebuilt. If bad blocks should be encountered during Rebuild, Add Drive, Host Write, or Regenerate Parity operation, the controller will first attempt to reconstruct affected data and those irrecoverable bad blocks are stated as bad and the controller return to host. Users have the option to abandon data on the unrecoverable sectors to continue rebuild in a degraded mode. Low quality drive handling comes with transparent resetting of hung hard drives.
Transparent reset of hung HDDs	Supported
Auto cache flush on critical conditions **(caching mode dynamic switch)**	When critical conditions occur, e.g., component failure, or BBU under charge, cached data will be flushed and the write policy will be changed to write-through mode. Configurable “Trigger Events” for Write-through/Write-Back Dynamic Switch.
RAID parity update tracking and recovery	Yes, to avoid write holes.
Host-side Ordered Tag support	Supports write commands with embedded Ordered Tags.
Drive identification (flash drive function)	Supported. Force a drive to light on its activity indicator for users to visually recognize its position in a configuration consisting of numerous disk drives.
Drive information listing	Supported. Drive vendor name, model number, firmware revision, capacity (blocks), serial number, narrow/wide and current sync. speed
Drive read/write test	Supported
Configuration on disks (Drive Roaming)	Will be supported in next release. The logical drive information is recorded on drive media. The logical drives can still be accessed if using different Galaxy RAID controllers/subsystems, e.g., drives removed and installed in a different subsystem.
Save/ restore NVRAM to / from disks	Supported. Save all the settings stored in the controller NVRAM to the logical drive members. Now this feature comes with an option whether to restore the previously saved password in case an administrator changed the password some time before or simply forgets the previous password.

Save / restore NVRAM to / from a file	Supported. Save all the settings stored in the controller NVRAM to a file (via GUI manager) on user's computer. Now this feature comes with an option whether to restore the previously saved password in case an administrator changed the password some time before.
Host-side 64-bit LBA support	Supports array configuration (logical drive, logical volume, or a partition of them) of a capacity up to 64TB.
Host LUN geometry: user configurable default geometry (Solaris OSes)	This feature comes with preset combinations of head, cylinder, and sector variables.
User configurable geometry range:	Sector: 32, 64, 127, 255 or Variable Head: 64, 127, 255 or Variable Cylinder: <1024, <32784,<65536 or Variable
Drive motor spin-up	Supported. The controller will send spin-up (start unit) command to each drive at the 4 sec. intervals.
Drive-side tagged command queuing	Supported. User adjustable up to 128 for each drive.
Host-side maximum queued I/O count	User adjustable up to 1024
Maximum concurrent host LUN connection	User adjustable up to 1024
Number of tags reserved for each Host-LUN connection	User adjustable up to 256
Controller shutdown	Flushes cached contents upon the detection of critical conditions, e.g., a high temperature condition persists for a long time.
Drive I/O timeout	User adjustable
I/O channel diagnostics	Supported; please contact your dealer for more details.
Power Saving	Idle and Spin-down modes
Maximum Drive Response Time (Guaranteed Latency I/O)	User adjustable from 160 to 960ms. If a disk drive fails to return data on read requests before the timeout value is exceeded, the array immediately generates data from the parity data and the other members of a logical drive.

17.3 A.3 Caching Operation

Write-back cache	Supported.
Write-through cache	Supported.

Supported memory type	DDR memory for enhanced performance. Fast Page Memory with Parity for enhanced data security.	
Read-ahead operation	Intelligent and dynamic read-ahead operation for processing sequential data requests.	
Multi-threaded operation	Yes, internal parameters adjusted in accordance with the number of outstanding I/Os.	
Scatter / Gather	Supported	
I/O sorting	Supported. Optimized I/O sorting for enhanced performance.	
Adaptive Write-back/Write-through switching	For a better performance when handling large sequential writes, firmware temporarily disables write-back cache and the synchronized cache operation between partner controllers if operating with dual-active RAID controllers. Firmware automatically restores the write-back mode when encountering random and small writes later.	
Periodic Cache Flush	Firmware can be configured to flush the cached contents in memory at every preset interval: If data integrity is of the concern, e.g., the lack of a battery backup protection. Cache flush on preset intervals to avoid the latency when cache memory is full due to write delays.	
Variable stripe size	RAID0	128
	RAID1	128
	RAID3	16
	RAID5	128
	RAID6	128
Caching Optimization		
Cache buffer sorting prior to cache flush operation.		
Gathering of writes during flush operation to minimize the number of I/Os required for parity update.		
Elevator sorting and gathering of drive I/Os.		
Multiple concurrent drive I/Os (tagged commands).		
Intelligent, predictive multi-threaded read-aheads.		
Multiple, concurrent host I/O threads (host command queuing).		

17.4 RAID Expansion

On-line RAID expansion	Supported. Capacity brought by array expansion is immediately ready for Host I/Os when its status changes from “EXPAND” to “INITIALIZING.” Initialization task is then completed in the background except when the logical array is stated as “INCOMPLETE” or “BAD;” e.g., has a failed member right after creation.
Mode-1 RAID expansion -add drive	Supported. Multiple drives can be added concurrently. Though not recommended, Add Drive can even be performed in the degraded mode.
Mode-2 RAID expansion – copy and replace drives	Supported. Replace members with drives of larger capacity.
Expand capacity with no extra drive bays required	Supported in Mode 2 RAID expansion, which provides “Copy and Replace Drive” function to replace drives with drives of greater capacity. Protect your investment for there is NO need for hardware upgrade, e.g., adding a new enclosure for the extra drives.
Operating system support for RAID expansion	No. No operating system driver required. No software needs to be installed for this purpose.

17.5 S.M.A.R.T. Support

Copy & replace drive	Supported. User can choose to clone a member drive showing symptoms of defects before it fails.
Drive S.M.A.R.T. support	Supported, with intelligent error handling implementations.
User selectable modes on the occurrence of S.M.A.R.T.-detected errors	Detect only Perpetual Clone: using a hot-spare to clone the drive reporting SMART errors; the hot-spare remains a clone drive Clone + Replace: using a hot-spare to replace the drive reporting SMART errors; the drive reporting errors is pulled offline Fail Drive: disband faulty drive from a logical drive.

17.6 Redundant Controller

Active-active redundant controller	Supported

Synchronized cache	Supported. Through one or multiple, dedicated synchronizing channels on a common backplane or external cabling. Synchronized cache over SCSI channels, Fibre loops, or SATA channels is supported. Synchronized cache can be disabled via a UI option when using write-through mode in a redundant controller configuration to prevent performance trade-offs.
Write-back cache enabled in redundant controller mode	Yes, with synchronized cache connection and mirrored cache between controllers.
Automatic failover	Yes (user's interaction necessary; e.g., to restart the software management console)
Automatic failback	Yes (user's interaction necessary)
Controller hot-swap	No need to shut down the failed controller before replacing the failed controller. Support online hot-swap of the failed controller. There is no need to reset or shutdown the failed controller. One controller can be pulled out during active I/Os to simulate the destructive controller failure.
Parity synchronization in redundant controller write-back mode to avoid write-hole	Supported.
No single-point-of-failure	Supported.
Automatic engagement of replacement controller	Supported.
Dynamic cache memory allocation	Yes. Cache memory is dynamically allocated, not fixed.
Environment management	Supported. SAF-TE, S.E.S., ISEMS (I2C interface), or S.E.S. over SAS links; and on-board controller voltage/temp monitor are all supported in both single and redundant controller mode. In the event of controller failure, services can be taken over by the surviving controller.
Cache Backup Module (CBM)	Supported. Battery backup modules support the transaction of cached data to flash memory on the occurrence of power outage. With EEPROM battery modules, firmware will be aware of the life expectancy of battery cells.

Load sharing	Supported. Workload can be flexibly divided between different controllers by assigning logical configurations of drives (LVs) to different RAID controllers.
User configurable channel mode	Supported. Channel modes configurable (SCSI or Fibre) as HOST or DRIVE on specific models .
Require a special firmware for redundant controller?	No.

17.7 Data Safety

Data Services	Snapshot, Volume Copy, Volume Mirror. Please refer to Galaxy Array Manager Manual for details.
Regenerate parity of logical drives	Supported. Can be manually executed to ensure that bad sectors do not cause data loss in the event of drive failure.
Scheduled Media Scan	Media Scan can be scheduled starting at a specified start time and repeated at regularly timed intervals. The start time and time intervals can be selected from drop-down menus. Start time is manually entered using its numeric representatives in the following order [MMDDhhmm[YYYY]], and it reads the date and time set for the controller's real-time clock. The selectable time intervals (the Execution Period) range from one (1) second to seven (7) weeks. Each such schedule can be defined to operate on individual hard drives, all members of a specified logical drive, or members of selected logical drives. Each schedule can include up to five (5) logical drives. The RS-232C terminal interface and RAIDWatch revision 2.0 support this functionality.
Bad block auto-reassignment	Supported. Automatic reassignment of bad block
Battery backup for cache memory	Supported. The battery backup unit supports cache memory when power failure occurs. The unwritten data in the cache memory can be committed to drive media when power is restored.
Verification on normal writes	Supported. Performs read-after-write during normal write processes to ensure data is properly written to drives.
Verification on rebuild writes	Supported. Performs read-after-write during rebuild write to ensure data is properly written to drives.

Verification on LD initialization writes	Supported. Performs read-after-write during logical drive initialization to ensure data is properly written to drives.
Drive S.M.A.R.T. support	Supported. Drive failure is predictable with reference to the different variables detected. Reaction schemes are selectable from Detect only, Perpetual Clone, Copy + Replace, and Fail Drive. These options help to improve MTBF.
Clone failing drive	Users may choose to clone data from a failing drive to a backup drive manually.
Automatic shutdown on over-temperature condition	Controller automatically enters an idle state (stops answering I/O requests) upon the detection of high-ambient temperature for an extended period of time.

17.8 System Security

Password protection	Supported. All configuration changes require the correct password (if set) to ensure system security. Password protection is also bundled with all user interfaces.
User-configurable password validation timeout	Supported. After certain time in absence of user interaction, the password will be requested again. This helps to avoid unauthorized operation when user is away.
SSL-enabled RAIDWatch Agents	Agents communicate to the controller through limited set of authorization options.

17.9 Environment Management

SAF-TE/S.E.S. support	Supported. The SAF-TE/S.E.S. modules can be connected to the drive channels. The RAID controller will detect errors from SAF-TE/S.E.S. devices or notify drive failures via SAF-TE/S.E.S. Both SAF-TE/S.E.S. via drive and device-self-interfaced methods are supported. Redundant SAF-TE/S.E.S. devices are supported Multiple S.E.S. devices are supported
Dynamic on-lining of enclosure services	Once an expansion unit (JBOD) with supported monitoring interface is combined with a RAID system, its status will be automatically polled.
SAF-TE/S.E.S. polling period	User configurable (50ms, 100ms, 200ms, 500ms, 1~60sec)
ISEMS (Galaxy Simple Enclosure	Supported via an I2C serial bus.

Management Service)	
Multiple SAF-TE/S.E.S. modules on the same channel	Supported.
Multiple SAF-TE /S.E.S. modules on different channels	Supported.
Mapping SAF-TE/S.E.S. device to host channel for use with host-based SAF-TE/S.E.S. monitoring	Supported.
Event Triggered Operation	When any of the following happens, the firmware disables write-back caching to minimize the chance of losing data: Battery, controller, cooling fan, or PSU failure The upper temperature thresholds are exceeded Low battery charge UPS AC loss or low battery charge The triggering factors are user-configurable
Multi-speed cooling fan control	Yes, firmware triggers high rotation speed in the event of elevated temperature or component failure, e.g., a fan failure.
Dual-LED drive status indicators	Supported. Both single-LED and dual-LED drive status indicators are supported.
SAF-TE/ S.E.S. temperature value display	Supported. Display the temperature value provided by enclosure SAF-TE/S.E.S. module (if available).
On-board controller voltage monitors	Supported. Monitors the 3.3V, 5V, and 12V voltage status. Event triggered thresholds user configurable.
On-board controller temperature sensors	Supported. Monitors the CPU and board temperature status. Event trigger threshold user configurable.
Enclosure redundant power supply status monitoring	Supported. SAF-TE/S.E.S./ISEMS
Enclosure fan status monitoring	Supported. SAF-TE/S.E.S/ISEMS
Enclosure UPS status monitoring	Supported. SAF-TE/S.E.S/ISEMS
Enclosure temperature monitoring	Supported. SAF-TE/S.E.S/ISEMS

17.10 User Interface

RAIDWatch on-board (Embedded RAIDWatch)	Out-of-band configuration and monitoring via Ethernet. http-based Embedded RAIDWatch interface that requires no

	installation efforts.
RS-232C terminal	Supports terminal modes: ANSI, VT-100, ANSI Color. Provides menu-driven user-friendly text-based, menu-driven interface.
Graphical user interface (Java-based GUI manager)	Provides user-friendly graphical interface. Communicates with RAID controller via Out-of-band Ethernet, In-band SCSI, In-band Fibre or SNMP traps.
SSH support	Secure Shell over Telnet supported
External interface API for customized host-based management	Supported.
LCD front panel	Provides easy access for user instinct operation.
Buzzer alarm	Warns users when any failures or critical events occur.

17.11 High Availability

Custom inquiry serial number	Custom Inquiry Serial Number (for support of multi-pathing software like Veritas, QLogic, etc).
Continuous rebuild	Rebuild automatically continues if power outage or operator errors occur during a rebuild.
Asymmetric Logical Unit Access (or later known as Target Port Group Service)	Support for multipath drivers to select an optimal I/O path and for more flexible utilization of internal I/O paths in the event of path failure or controller failover/failback.
High Availability hardware modules	Transparent controller failover/failback. IP address of the 10/100BaseT Ethernet port is handed over to a surviving controller in the event of a single controller failure. Multiple drive channel across the backplane to disk drives.

18 System Functions:

Upgrading Firmware

18.1 Upgrading Firmware

The RAID controller’s firmware resides in flash memory that can be updated through the COM port, LAN port, or via In-band SCSI/Fibre. New releases of firmware, when available, can be emailed to you. The file available is usually a self-extracting file that contains the following:

- FW30Dxyz: Firmware Binary (where "xyz" refers to the firmware version)
- B30Buvw: Boot Record Binary (where "uvw" refers to the boot record version)
- README.TXT: Read this file first before upgrading the firmware/boot record. It contains the most up-to-date information which is very important to the firmware upgrade and usage.

These files must be extracted from the compressed file and copied to a directory in boot drive.

18.1.1 Sample Upgrade Flowchart

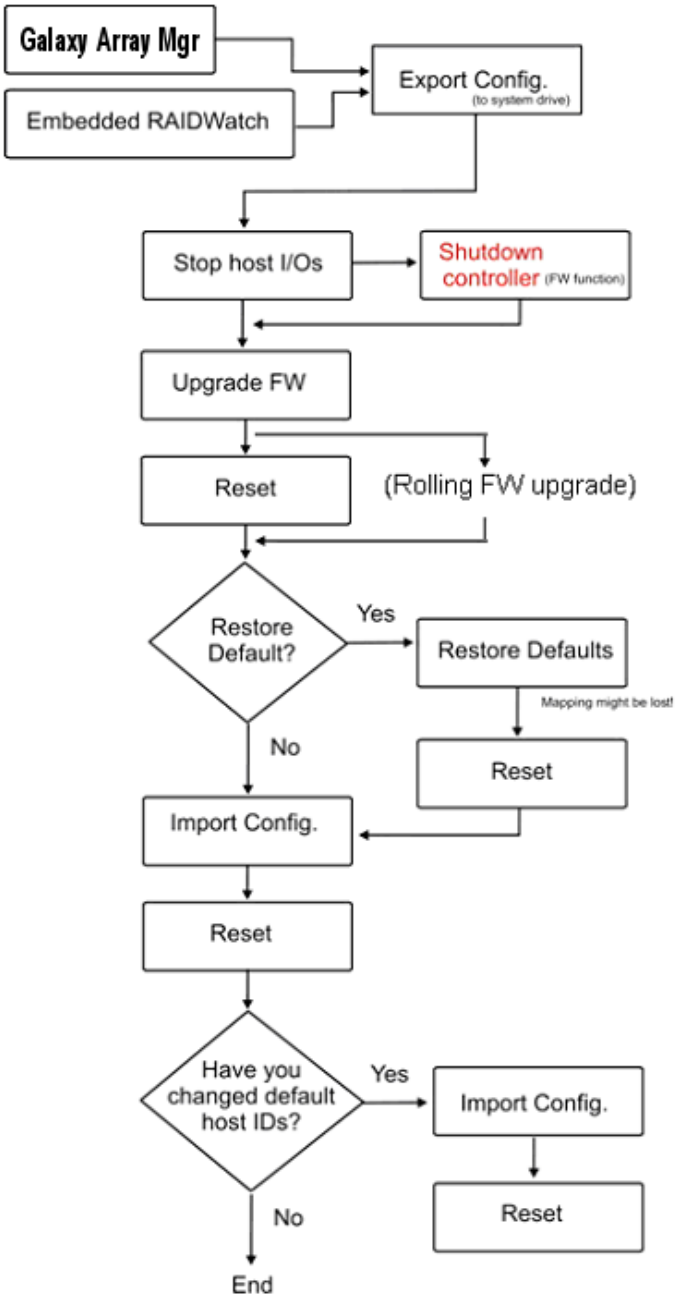

18.1.2 Note for Redundant Controller Firmware Upgrade

- Host I/Os will not be interrupted during the download process. After the download process is completed, user should find a chance to reset the controller for the new firmware to take effect.
- A controller used to replace a failed unit in a dual-controller system is often running a newer release of firmware version. To solve the contention, make sure the firmware on a replacement controller is downgraded to that running on the

surviving controller.

- Allow the downloading process to finish. Do not reset or turn off the computer or the controller while it is downloading the file. Doing so may result in an irrecoverable error that requires the controllers needing service.
- When upgrading the firmware, check the boot record version that comes with it. If the boot record version is different from the one installed on the surviving controller previously, the new boot record binary must be installed.
- Restore Default might be necessary when migrating firmware between major revisions.. Migration across many revisions may not be supported due to the differences in system hardware. Restore Default can erase the existing LUN mappings. Please consult technical support if you need to apply a very new firmware.
- Saving NVRAM (firmware configuration) to a system drive preserves all configuration details including host LUN mappings.
- Whenever host channel IDs are added or removed, you need to reset the system for the configuration to take effect. That is why you have to import your previous configuration and reset again to bring back the host LUN mappings if you have host IDs different from system defaults.

18.2 Upgrading Firmware Using Galaxy Array Manager

To upgrade the subsystem firmware, you need to work on the Configuration Manager in Galaxy Array Manager.

1. Obtain a new firmware file from your dealer.
2. Store the firmware file and the optional boot record file in a local directory. If you need to obtain an updated version of firmware and boot record (to use a new feature, to fix a bug, etc.) contact technical support.
3. Select the In-Band or Out-of-Band group in the sidebar of Galaxy Array Manager Commander and click the Configuration Manager icon or select the *Tools* > *Configuration Manager* menu.

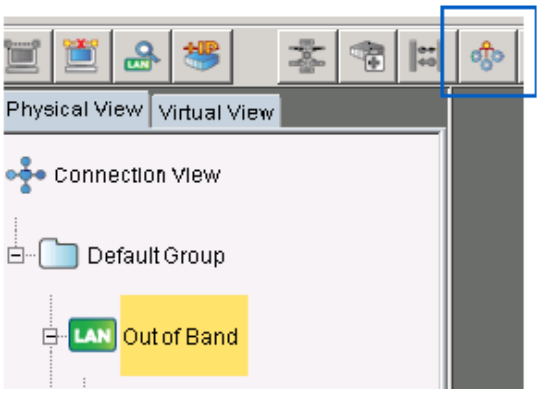

4. Select the storage systems you want to upgrade and click the Connect button.

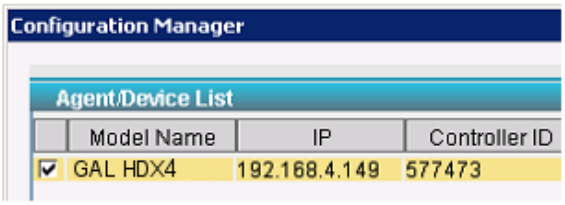

NOTE

You can select multiple storage subsystems and upgrade their firmware at once.

5. In the Configuration Manager window that appears, select the Maintenance tab. Select the type of firmware upgrade and click the Apply button.

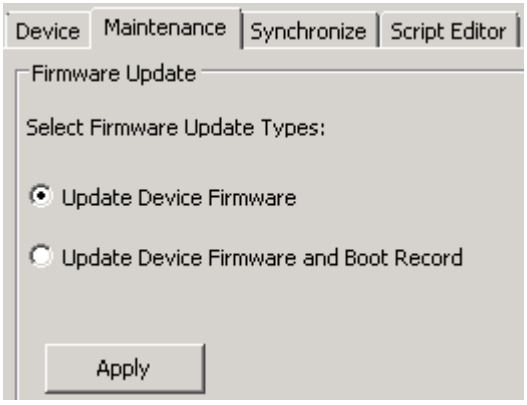

Update Device Firmware: Updates the storage system's firmware with the latest version.

Update Device Firmware and Boot Record: Updates the storage system's firmware and boot record with the latest version.

6. Select the firmware file (and boot record file) in the local folder and continue the process following the instructions.

18.3 Upgrading Firmware Using RS-232C Terminal Emulation

The firmware can be downloaded to the Galaxy RAID controller/subsystem by using an ANSI/VT-100 compatible terminal emulation program. Whichever terminal emulation program is used must support the ZMODEM file transfer protocol. The following example uses the HyperTerminal in Windows NT®. Other terminal emulation programs (e.g., Telix and PROCOMM Plus) can perform the firmware upgrade as well.

18.3.1 Upgrading Both Boot Record and Firmware Binaries

Go to: System Functions > Controller Maintenance > Update Firmware

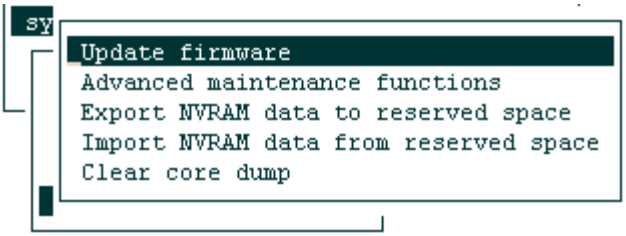

1. You may see the message “Recommended baud rate for update firmware while processing I/O is 9600.” Select Yes.
2. Set ZMODEM as the file transfer protocol of your terminal emulation software.
3. Send the Boot Record Binary to the controller. In HyperTerminal, go to the "Transfer" menu and choose "Send file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).
4. After the Boot Record has been downloaded, send the Firmware Binary to the controller. In HyperTerminal, go to the "Transfer" menu and choose "Send file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).
5. When the Firmware completes downloading, the controller will automatically reset itself. For a newer version of firmware, you need to manually reset the subsystem/controller for the new firmware to take effect.

19 Appendix

19.4 Default System Settings

Event Triggered Operations:

Controller failure	Disabled
BBU low or failed	Enabled
UPS AC power loss	Disabled
Power supply failure	Disabled
Fan failure	Disabled
Temperature exceeds threshold	Disabled

Host-side Parameters

Maximum Queued IO Count	1024
LUNs per host ID	8
Max. Number of Concurrent Host-LUN Connections	4
Number of Tags Reserved for Each Host-LUN Connection	32
Fibre Connection option	Loop only

Peripheral Device Parameters (for in-band management access only)

Peripheral device type	Enclosure Service Device (0xD)
Peripheral device qualifier	Connected
Device support removable media	Disabled
LUN applicability	First Undefined LUN (automatically selected by firmware)

Cylinder/Head/Sector- variables	N/A
Drive-side Parameters	
Disk Access Delay Time	Per product interface; will be a larger values in multi-enclosure applications
Drive I/O Timeout	7 seconds
Max. Tag Count	8: Fibre Channel
Periodic SAF-TE and SES Check Time	30 seconds
Auto Rebuild on Drive Swap check time	15 seconds
Drive Predictable Failure Mode (S.M.A.R.T.)	Disabled
Drive Delayed Write	Enabled (single-controller, w/o BBU) Disabled (dual-controller, w/ BBU)
Drive Spindown Idle Delay	Disabled
Drive-side Parameters	
Disk Access Delay Time	Per product interface; will be a larger values in multi-enclosure applications
Drive I/O Timeout	7 seconds
Max. Tag Count	8: Fibre Channel
Periodic SAF-TE and SES Check Time	30 seconds
Auto Rebuild on Drive Swap check time	15 seconds
Drive Predictable Failure Mode (S.M.A.R.T.)	Disabled
Drive Delayed Write	Enabled (single-controller, w/o BBU) Disabled (dual-controller, w/ BBU)

Drive Spindown Idle Delay	Disabled

Voltage & Temperature Parameters

+3.3V thresholds	3.6V – 2.9V
+5V thresholds	5.5V – 4.5V
+12V thresholds	13.2V - 10.8V
CPU temperature	90 - 5°C
Board temperature (RAID controller board)	80 - 5°C

The thresholds for other sensors within the chassis are not user-configurable. It is user's responsibility to maintain a reasonable ambient temperature, e.g., below 35°C, and stable power source at the installation site.

Disk Array Parameters

Rebuild Priority	Normal
Verification on Write	
Verification on LD Initialization	Disabled
Verification on LD Rebuild	Disabled
Verification on Normal Drive Writes	Disabled
Max. Drive Response Timeout	Disabled

19.5 ASCII Code Table

(Supported Characters for controller name, password, WWPN port nick names, etc.)

Note that ox5c back slash is not supported.

032	040	020	00100000	SP	(Space)
033	041	021	00100001	!	(exclamation mark)
034	042	022	00100010	"	(double quote)
035	043	023	00100011	#	(number sign)
036	044	024	00100100	$	(dollar sign)

037	045	025	00100101	%	(percent)
038	046	026	00100110	&	(ampersand)
039	047	027	00100111	'	(single quote)
040	050	028	00101000	(	(left/open parenthesis)
041	051	029	00101001	)	(right/closing parenth.)
042	052	02A	00101010	*	(asterisk)
043	053	02B	00101011	+	(plus)
044	054	02C	00101100	,	(comma)
045	055	02D	00101101	-	(minus or dash)
046	056	02E	00101110	.	(dot)
047	057	02F	00101111	/	(forward slash)
048	060	030	00110000	0	
049	061	031	00110001	1	
050	062	032	00110010	2	
051	063	033	00110011	3	
052	064	034	00110100	4	
053	065	035	00110101	5	
054	066	036	00110110	6	
055	067	037	00110111	7	
056	070	038	00111000	8	
057	071	039	00111001	9	
058	072	03A	00111010	:	(colon)
059	073	03B	00111011	;	(semi-colon)
060	074	03C	00111100	<	(less than)
061	075	03D	00111101	=	(equal sign)
062	076	03E	00111110	>	(greater than)

063	077	03F	00111111	?	(question mark)
064	100	040	01000000	@	(AT symbol)
065	101	041	01000001	A	
066	102	042	01000010	B	
067	103	043	01000011	C	
068	104	044	01000100	D	
069	105	045	01000101	E	
070	106	046	01000110	F	
071	107	047	01000111	G	
072	110	048	01001000	H	
073	111	049	01001001	I	
074	112	04A	01001010	J	
075	113	04B	01001011	K	
076	114	04C	01001100	L	
077	115	04D	01001101	M	
078	116	04E	01001110	N	
079	117	04F	01001111	O	
080	120	050	01010000	P	
081	121	051	01010001	Q	
082	122	052	01010010	R	
083	123	053	01010011	S	
084	124	054	01010100	T	
085	125	055	01010101	U	
086	126	056	01010110	V	
087	127	057	01010111	W	
088	130	058	01011000	X	

089	131	059	01011001	Y	
090	132	05A	01011010	Z	
091	133	05B	01011011	[	(left/opening bracket)
092	134	05C	01011100	\	(back slash)
093	135	05D	01011101	]	(right/closing bracket)
094	136	05E	01011110	^	(caret/circumflex)
095	137	05F	01011111	_	(underscore)
096	140	060	01100000	`	
097	141	061	01100001	a	
098	142	062	01100010	b	
099	143	063	01100011	c	
100	144	064	01100100	d	
101	145	065	01100101	e	
102	146	066	01100110	f	
103	147	067	01100111	g	
104	150	068	01101000	h	
105	151	069	01101001	i	
106	152	06A	01101010	j	
107	153	06B	01101011	k	
108	154	06C	01101100	l	
109	155	06D	01101101	m	
110	156	06E	01101110	n	
111	157	06F	01101111	o	
112	160	070	01110000	p	
113	161	071	01110001	q	
114	162	072	01110010	r	

115	163	073	01110011	s	
116	164	074	01110100	t	
117	165	075	01110101	u	
118	166	076	01110110	v	
119	167	077	01110111	w	
120	170	078	01111000	x	
121	171	079	01111001	y	
122	172	07A	01111010	z	
123	173	07B	01111011	{	(left/opening brace)
124	174	07C	01111100	|	(vertical bar)
125	175	07D	01111101	}	(right/closing brace)

RORKE DATA®
/ An Avnet Company /

Rorke Data, An Avnet Company
7626 Golden Triangle Drive, Eden Prairie, MN 55344, USA
» Toll Free 1.800.328.8147 » Phone 1.952.829.0300 » Fax 1.952.829.0988